arXiv:1402.0815v2 [math.GT] 5 Feb 2014

Embeddability in the 3-sphere is decidable

Jiří Matoušek[1,2,a,b], Eric Sedgwick[3,a], Martin Tancer[1,4,5,a], and Uli Wagner[4,a,c]

[1]Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic.
[2]Department of Computer Science, ETH Zürich, 8092 Zürich, Switzerland
[3]School of Computing, DePaul University, 243 S. Wabash Ave, Chicago, IL 60604, USA
[4]IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria
[5]Part of this work was done when the third author was affiliated with Institutionen för matematik, Kungliga Tekniska Högskolan, Linstedsvägen 25, 100 44 Stockholm
[a]Supported by the ERC Advanced Grant No. 267165.
[b]Partially supported by Grant GRADR Eurogiga GIG/11/E023.
[c]Supported by the Swiss National Science Foundation (Grants SNSF-PP00P2-138948 and SNSF-200020-138230).

Abstract

We show that the following algorithmic problem is decidable: given a 2-dimensional simplicial complex, can it be embedded (topologically, or equivalently, piecewise linearly) in $\mathbb{R}^3$? By a known reduction, it suffices to decide the embeddability of a given triangulated 3-manifold X into the 3-sphere S^3. The main step, which allows us to simplify X and recurse, is in proving that if X can be embedded in S^3, then there is also an embedding in which X has a *short meridian*, i.e., an essential curve in the boundary of X bounding a disk in $S^3 \setminus X$ with length bounded by a computable function of the number of tetrahedra of X.

1 Introduction

The embeddability problem. Let $\mathrm{EMBED}_{k\to d}$ be the following algorithmic problem: given a finite simplicial complex K of dimension at most k, does there exist a (piecewise linear) embedding of K into $\mathbb{R}^d$?

A systematic investigation of the computational complexity of this problem was initiated in [MTW11]; earlier it was known that $\mathrm{EMBED}_{1\to 2}$ (graph planarity) is solvable in linear time, so is $\mathrm{EMBED}_{2\to 2}$ [GR79], and for every $k \geq 3$ fixed, $\mathrm{EMBED}_{k\to 2k}$ can be decided in polynomial time (this is based on the work of Van Kampen, Wu, and Shapiro; see [MTW11]).

For dimension $d \geq 4$, there is now a reasonably good understanding of the computational complexity of $\mathrm{EMBED}_{k\to d}$: for all k with $(2d-2)/3 \leq k \leq d$ it is NP-hard (and even undecidable if $k \geq d-1 \geq 4$) [MTW11], while for $k < (2d-2)/3$ it is polynomial-time solvable, assuming d fixed, as was shown in a series of papers on computational homotopy theory [ČKM+13, ČKM+12, KMS13, ČKV13]. (However, the cases with $(2d-2)/3 \leq k$ known to be NP-hard but not proved undecidable are still intriguing.)

Thus, the most significant gap up until now has been the cases $d = 3$ and $k = 2, 3$, and in particular, after graph planarity ($\text{EMBED}_{1\to 2}$), the problem $\text{EMBED}_{2\to 3}$ can be regarded as the most intuitive and probably practically most relevant case.

Embeddability in $\mathbb{R}^3$. Here we close this gap, at least as far as decidability is concerned.

Theorem 1.1. *The problem* $\text{EMBED}_{2\to 3}$ *is algorithmically decidable. That is, there is an algorithm that, given a* 2*-dimensional simplicial complex* K*, decides whether* K *can be embedded (piecewise linearly, or equivalently, topologically) in* $\mathbb{R}^3$*.*

Let us remark that one can naturally consider (at least) three different kinds of embeddings of a simplicial complex K in $\mathbb{R}^d$, illustrated in the next picture for a 1-dimensional complex (graph):

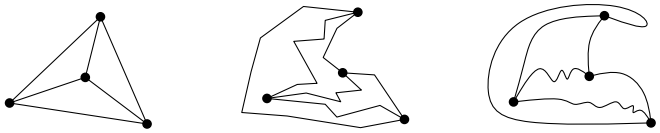

linear | piecewise linear (PL) | topological

For *linear embeddings*, also referred to as *geometric realizations*, each simplex of K should be mapped affinely to a (straight) geometric simplex in $\mathbb{R}^d$. This kind of embeddability is decidable in PSPACE regardless of the dimensions, and it is *not* what we consider here.

For *piecewise linear*, or *PL*, embeddings, one seeks a linear embedding of some (arbitrarily fine) subdivision of K. Finally, for a *topological embedding*, K is embedded by an arbitrary injective continuous map.

While topological and PL embeddability need not coincide for some ranges of dimensions, for ambient dimension $d = 3$, they do,[1] and this is the notion of embeddability considered here.

An algorithm for $\text{EMBED}_{3\to 3}$ can be obtained from Theorem 1.1 by a simple reduction, given in Section 12.

Corollary 1.2. *The problem* $\text{EMBED}_{3\to 3}$ *is decidable as well.*

Thickening to 3-manifolds. For a 2-complex K, (PL) embeddability in $\mathbb{R}^3$ is easily seen to be equivalent to embeddability in S^3, and from now on, we work with S^3 as the target.

The first step in our proof of Theorem 1.1 is testing whether a given simplicial 2-complex K embeds in any 3-dimensional manifold at all.

Let us suppose that there is an embedding $f\colon K \to M$ for some 3-manifold M (without boundary), and take a sufficiently small closed neighborhood X of the image $f(K)$ in M—the technical term here is a *regular neighborhood*. Then X is a 3-manifold with boundary, called a 3-*thickening* of K.

There is an algorithm, due to Neuwirth [Neu68] (see also [Sko95] for an exposition) that, given K, tests whether it has any 3-thickening, and if yes, produces a finite list of all possible 3-thickenings, up to homeomorphism, as triangulated 3-manifolds with boundary (*without* the knowledge of M). Then K embeds in S^3 iff one of its 3-thickenings does. Hence it suffices to prove the following.

[1] For complexes of dimension $k = 2$, this follows from[Bin59, Pap43], see also [MTW11] for more details and references; for complexes of dimension $k = 3$, this follows from the reduction in Section 12.

Theorem 1.3. *There is an algorithm that, given a triangulated 3-manifold X with boundary, decides whether X can be embedded in S^3.*

Concerning the running time. Our proof does provide an explicit running time bound for the algorithm, but currently a rather high one, certainly primitive recursive but even larger than an iterated exponential tower. Thus, we prefer to keep the bounds unspecified, in the interest of simplicity of the presentation.

By refining our techniques, it might be possible to show the problem to lie in the class NP. Going beyond that may be quite challenging: indeed, as observed in [MTW11], $\mathrm{EMBED}_{2\to 3}$ is at least as hard as the problem of recognizing S^3 (that is, given a simplicial complex, decide whether it is homeomorphic to S^3). The latter problem is in NP [Iva08, Sch04], but nothing more seems to be known about its computational complexity (e.g., polynomiality or NP-completeness).

Related work. There is a vast amount of literature on computational problems for 3-manifolds and knots. Here we give just a sample; further background and references can be found in the sources cited below and in [AHT06]. A classical result is Haken's algorithm deciding whether a given polygonal knot in $\mathbb{R}^3$ is trivial [Hak61]. More recently, this problem was shown to lie in NP [HLP99], and, assuming the Generalized Riemann Hypothesis, in coNP as well [Kup11]. The knot equivalence problem is also decidable [Hak61, Hem79, Mat97], but nothing seems to be known about its complexity status.

Closer in spirit to the problem investigated here are algorithms for deciding whether a given 3-manifold is homeomorphic to S^3, already mentioned above [Rub95, Tho94, Iva08, Sch04].

An important special case of Theorem 1.3 is testing embeddability into S^3 for an X whose boundary is a single torus; this amounts to recognizing knot complements and was solved in [JS03]. Some of the ideas in that work are used in our proof, but most of the argument is fairly different.

In a different direction, Tonkonog [Ton11] provided an algorithm for deciding whether a given 3-manifold X with boundary embeds into *some* homology 3-sphere[2] (which may depend on X). His methods are completely different from ours (except for using a 3-thickening to pass from 2-dimensional complexes to 3-manifolds), and it seems to be only loosely related to the problems investigated here.

Future directions. Besides the obvious questions of finding a more efficient algorithm, say one in NP, and/or proving hardness results, one may consider embeddability into other 3-manifolds M besides S^3. We believe that this may be within reach of the methods used here, but definitely a number of issues would have to be settled.

The main technical contribution. Our algorithm relies on a large body of work in 3-dimensional topology.

When we talk about a *surface* in X, unless explicitly stated otherwise, we always mean a 2-dimensional manifold F with boundary *properly embedded* in X, that is, with $\partial F \subset \partial X$. Similarly, curves are considered properly embedded in a surface, so a connected curve can be a loop in the interior of the surface or an arc connecting two points of the boundary. Two properly embedded surfaces F and F' are *isotopic* if they are embeddings of the same surface F_0 and there is a continuous family of proper embeddings $F_0 \to X$ starting with F and ending

[2] A 3-manifold whose homology groups are the same as those of S^3.

with F'. An similar definition of isotopy applies to curves embedded in surfaces.

As in almost all algorithms working with 3-manifolds, we use Haken's method of *normal curves and surfaces*, actually in a slightly extended form. Here we recall them very briefly; we refer to [Hem76, JT95] for background, and in Section 5 below we will discuss a variant.

A *normal curve* in a triangulated 2-dimensional surface F intersects every triangle in finitely many disjoint pieces, which we can think of as straight segments, as in the left picture:

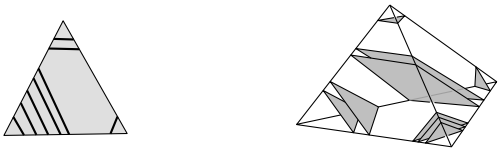

The main point is that such a curve is described, up to isotopy, purely combinatorially: namely, for every triangle T, there are just three types of segments of the curve inside, and it is enough to specify the number of segments for each type, for each T. In the picture, the numbers are $5, 2, 1$.

Similarly, a *normal surface* in a triangulated 3-manifold intersects each tetrahedron in finitely many of disjoint pieces, each of them a triangle or a quadrilateral, as in the right picture above. This time there are seven types of pieces, four triangular and three quadrilateral, per tetrahedron (although no two types of quadrilateral pieces may coexist in a single tetrahedron, since they would have to intersect, which is not allowed). So a normal surface F in a 3-manifold with t tetrahedra can be described by a vector of $7t$ nonnegative integers. This vector is called the *normal vector* of F.

A *normal isotopy* is an isotopy during which the intermediate curve or surface stays normal; in particular, it may not cross any vertex of the triangulation.

Going back to embeddings, we first simplify the situation using a result of Fox [Fox48], which allows us to assume that the complement of the supposed embedding of X in S^3 is a union of handlebodies.[3] (These handlebodies may be knotted or linked in S^3, though, as in the picture at the beginning of the next section.) This assumption is quite important and nontrivial; for example, we note that if X is a solid torus, it can also be embedded in S^3 in a knotted way, so that the complement is not homeomorphic to a solid torus.

Thus, now we ask if there is a way of "filling" each component of ∂X with a handlebody so that the resulting closed manifold is homeomorphic to S^3. Spherical boundary components are easy, since there is only one way, up to homeomorphism, of filling a spherical boundary component with a ball. However, already for a toroidal component there are infinitely many nonequivalent ways of filling it with a solid torus. Indeed, the filling can be done in such a way that a circle α on the toroidal component of ∂X, as in the left picture,

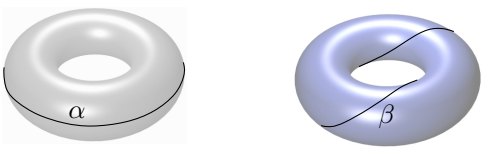

is identified with a curve β on the boundary of the solid torus, shown in the right picture, where β may wind around the solid torus as many times as desired. For boundary components

[3] A *handlebody* is a ball with (solid) handles, or equivalently, a 3-thickening of a finite connected 1-dimensional complex (graph).

of higher genus, there are also infinitely many ways of filling, and their description is still more complicated. For every specific way of filling the boundary components of X with handlebodies we could test whether the resulting closed manifold is an S^3, but we cannot test all of the infinitely many possibilities. This is the main difficulty we have to overcome to get an algorithm.

Next, by more or less standard considerations, we can make sure that there is no "way of simplifying X by cutting along a sphere or disk"—in technical terms, we may assume that X is *irreducible*, that is, every 2-sphere embedded in X bounds a ball in X, and that X has an *incompressible boundary*, i.e., any curve in ∂X bounding a disk in X also bounds a disk in ∂X.

For dealing with such an X, the following result is the key:

Theorem 1.4. *Let X be an irreducible 3-manifold, neither a ball nor an S^3, with incompressible boundary and with a 0-efficient triangulation $\mathcal{T}$. If X embeds in S^3, then there is also an embedding for which X has a* short meridian *γ, i.e., an essential*[4] *normal curve $\gamma \subset \partial X$ bounding a disk in $S^3 \setminus X$ such that the length of γ, measured as the number of intersections of γ with the edges of $\mathcal{T}$, is bounded by a computable function of the number of tetrahedra in $\mathcal{T}$.*

In this theorem, 0-*efficient triangulation* is a technical term introduced in [JR03], whose definition will be recalled later in Section 7. We are using 0-efficient triangulations in order to exclude non-trivial normal disks and 2-spheres in X.

We should also mention that the triangulations commonly used in 3-dimensional topology, and also here, are not simplicial complexes in the usual sense—they are still made by gluing (finitely many) tetrahedra by their faces, but any set of gluings that produces a manifold is allowed, even those that identify faces of the same tetrahedron. As a result, a particular tetrahedron may not have four distinct faces, six distinct edges and four distinct vertices. In particular, 0-efficient triangulations of the manifolds we consider have a single vertex in each boundary component and none in the interior, all edges in the boundary form loops. This is the necessary result of modifying a triangulation by collapsing simplices, a triangular face to an edge or to a vertex, etc.; see [JR03, Sec. 2.1] for a thorough discussion. There is even a mind-boggling one-tetrahedron one-vertex triangulation of the solid torus obtained by gluing a pair of faces of a single tetrahedron, see [JS03].

Let us remark that X as in the theorem need not have a short meridian for every possible embedding, even if we assume that the complement consists of handlebodies. For example, if X is a thickened torus (a torus times an interval), we can embed it so that the curves bounding disks in $S^3 \setminus X$ are arbitrarily long w.r.t. a given triangulation of X. We must sometimes change the embedding to get short meridians.

It is also worth mentioning that this problem does not occur if ∂X is a single torus, i.e., the knot complement case. Here a celebrated theorem of Gordon and Luecke [GL89] makes sure that there is only one embedding, up to a self-homeomorphism of S^3, and the meridian is unique up to isotopy. This is why the single-torus boundary case solved in [JS03] is significantly easier than the general case.

[4]Meaning that γ does not bound a disk in ∂X.

2 An outline of the arguments

Our algorithm for Theorem 1.3, deciding the embeddability of a given 3-manifold X in S^3, for the case of X irreducible and with incompressible boundary, consists in testing every possible normal curve $\gamma \subset \partial X$ of length bounded as in Theorem 1.4. For each such candidate γ, we construct a new manifold $X' = X'(\gamma)$ by *adding a 2-handle* to X along γ, which means that we glue a disk bounded by γ to the outside of X and thicken it slightly, as illustrated in the following picture:

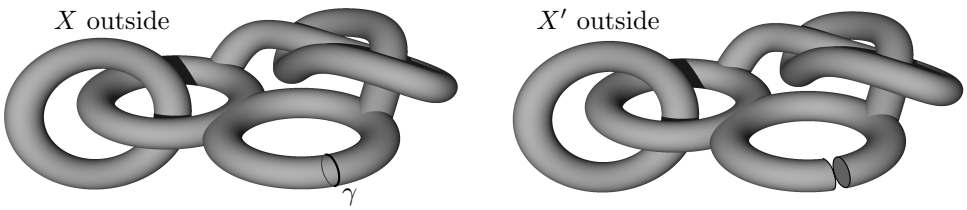

Here X is the complement of the union of two (linked) handlebodies, a knotted solid 3-torus and a solid torus, and for X', the solid 3-torus in the complement has been changed to a solid double torus.

Then we test the embeddability of each $X'(\gamma)$ recursively, and X is embeddable iff at least one of the $X'(\gamma)$ is. It is not hard to show that the algorithm terminates, using the vector of genera of the boundary components of X; see Section 3.

The proof of Theorem 1.4 occupies most of the paper and has many technical steps. In this section we give an outline.

We assume X to be embedded in S^3, the complement being a union of handlebodies, and we apply a result of Li [Li10] stating that there is a planar surface (i.e., a disk with holes) $P \subset X$ that is "stuck" in its position in a suitable sense (namely, P is either essential,[5] or strongly irreducible and boundary strongly irreducible) and is meridional or almost meridional.

Here an essential curve $\gamma \subset \partial X$ is a *meridian* in a given embedding of X in S^3 if it bounds a disk in $S^3 \setminus X$. The surface P is *meridional* if each component of ∂P is a meridian, and it is *almost meridional* if all components of ∂P but one are meridians. (Actually, Li has yet another case in his statement, but as we will check, that case can be reduced to the ones given above; see Lemma 4.4.) The next picture illustrates a meridional P in the case where X is embedded in S^3 as the complement of a solid torus neighborhood of the figure '8' knot:

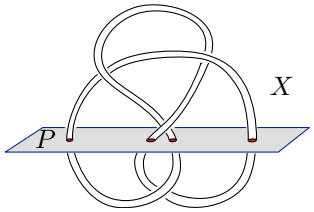

Next, by choosing P as above with suitable minimality properties, one can make sure that P is normal or almost normal[6] for the given triangulation. For the case of P essential, this

[5]The precise definitions of essential, strongly irreducible, and boundary strongly irreducible are somewhat complicated and we postpone them to Section 4.

[6]An *almost normal* surface is like a normal surface except that in at most one tetrahedron we also allow, in addition to the triangular and quadrangular pieces, one of two types of *exceptional pieces*, namely, a *tube* or an *octagon*; see Section 5.

is an old result going back to Haken and Schubert (and for our notion of complexity of P, a proof is given in Section 7), while for P strongly irreducible and boundary strongly irreducible this follows from [BDTS12]; also see [Sto00] for the case of a strongly irreducible surface in a closed manifold. It remains to show that, in this setting, at least one of the meridians in ∂P must be short.

Here we apply an *average length estimate*, which is an idea of Jaco and Rubinstein appearing in [JS03, JRS09].

Let $\gamma_1, \ldots, \gamma_b$ be the components of ∂P, and let $\ell(\partial P) = \sum_{i=1}^{b} \ell(\gamma_i)$ be the boundary length of P. We know that all the γ_i but at most one are meridians. The length of the shortest meridian is bounded by the average $\ell(P)/(b-1)$, and we want to bound this average by a (computable) function of t, the number of tetrahedra in the triangulation $\mathcal{T}$ of X.

Now by the theory of normal surfaces, the (almost) normal surface P can be written as a normal sum[7] of fundamental surfaces in X,

$$P = \sum_i k_i F_i, \tag{1}$$

where the k_i are positive integers and the F_i are surfaces from a finite collection; their number, as well as $\ell_{\max} := \max_i \ell(\partial F_i)$ can be bounded by a (computable) function of t alone, and does not depend on P.

Since the boundary length is additive w.r.t. normal sum, we have $\ell(\partial P) = \sum_i k_i \ell(\partial F_i) \leq \ell_{\max} K$, where $K := \sum_i k_i$ is the number of fundamental summands in the expression for P, and so it suffices to show that $K \leq Cb$, with some computable function $C = C(t)$.

The basic version of the average-length estimate uses the Euler characteristic χ as an accounting device. Since χ is additive as well, $\chi(P) = \sum_i k_i \chi(F_i)$. Since P is a planar surface with b boundary components, we have $\chi(P) = 2 - b$.

Now an ideal situation for the average-length estimate (which we cannot guarantee in our setting) is when $\chi(F_i) \leq -1$ for every i; in other words, none of the summands is a disk, 2-sphere, annulus, Möbius band, or torus (or projective plane or Klein bottle, but these cannot occur in X embedded in S^3). Then we get $b - 2 = -\chi(P) = \sum_i k_i(-\chi(F_i)) \geq K$, and we are done (even with $C = 1$).

In our actual setting, the summands with $\chi > 0$, i.e., spheres and disks, are excluded by the 0-efficient triangulation of X. We also need not worry about torus summands, since they have empty boundary and thus do not contribute to $\ell(\partial P)$. The real problem are annuli (and Möbius bands, but since twice a Möbius band, in the sense of normal sum, is an annulus, Möbius bands can be handled easily once we deal with annuli).

There are two kinds of annuli, which need very different treatment: the essential ones, and the *boundary parallel* ones. Here an annulus $A \subset X$ is *boundary parallel* if it can be isotoped to an annulus $A' \subset \partial X$ with $\partial A' = \partial A$ while keeping the annulus boundary fixed. Boundary parallel annuli do not occur for P essential, but they might occur for the case of P strongly irreducible and boundary strongly irreducible.

To deal with the annulus summands, we first construct what we call an *annulus curve* $\alpha \subset \partial X$. This is the boundary of a maximal collection $\mathcal{A}$ of essential annuli, maximal in the sense that each of the two boundary curves of every other essential annulus, after a suitable

[7]For normal surfaces F, F_1, F_2 in a triangulation $\mathcal{T}$ of X, F is called the *normal sum* of F_1 and F_2 if $\vec{v}(F) = \vec{v}(F_1) + \vec{v}(F_2)$, where $\vec{v}(F)$ denotes the normal vector of F. Similarly for almost normal surfaces, where we have extra coordinates in $\vec{v}(F)$ for the exceptional types of pieces; in this case, at least one of F_1, F_2 has to be normal. Also see Section 5.

normalization, either intersects α or is normally isotopic to a component of α. We bound the length of α by a computable function of t, and $|\alpha \cap P|$, the number intersections of α with P, by $C'b$, for some computable $C' = C'(t)$, again assuming P minimal in a suitable sense. For obtaining this bound we may need to change the embedding of X, and we also use results about "untangling" a system of curves on a surface by a boundary-fixing self-homeomorphism from [MSTW13].

Similarly, we construct a collection Γ of curves that helps to deal with boundary parallel annuli: those that have minimal boundary in a suitable sense either intersect α, or their boundaries are normally isotopic to components of α or curves from Γ.

Having constructed such an α and Γ, we work with normal curves and surfaces in a "marked" sense, which also takes into account the position of the curves and surfaces w.r.t. α and Γ. This, in particular, makes the number of intersections with α additive w.r.t. the marked normal sum, which in turn allows us to bound the number of annulus summands in (1), both boundary parallel and essential, that intersect α by $C'b$.

Then we might have boundary-parallel annulus summands that avoid α, but we show that those do not occur at all, since they would contradict the minimality of P.

Finally, there remain essential annuli that have a boundary component parallel to a component of α. Here we show that if such an annulus had the coefficient k_i in (1) at least $|\alpha \cap \partial P| \leq C'b$, then there is a self-homeomorphism of X, namely, a *Dehn twist* in the annulus, that makes P simpler, contradicting its supposed minimality. (Here we may again modify the assumed embedding of X in S^3 in order to get a short meridian—and, as we have remarked, some such modification is necessary in the proof, since some embeddings may not have short meridians.) Hence for these essential annuli, too, the coefficients are bounded by a linear function of b. This concludes the proof.

3 The algorithm

If X embeds in S^3, then it is orientable, and orientability can easily be tested algorithmically (e.g., by a search in the dual graph of the triangulation, or by computing the relative homology group $H_3(X, \partial X)$). So from now on, we assume X orientable. In this situation, the boundary of X is a compact orientable 2-manifold, and thus each component is a 2-sphere with handles.

We describe a recursive procedure `EMB`(X) that accepts a triangulated orientable 3-manifold with boundary and returns TRUE or FALSE depending on the embeddability of X in S^3. (With some more effort, for the TRUE case, we could also recover a particular embedding, but we prefer simplicity of presentation.) The procedure works as follows.

1. (Each component separately) Let $X_1, \ldots, X_k$ be the connected components of X. If $k > 1$, test if there is an S^3 among the X_i (several algorithms are available for that [Rub95, Tho94, Iva08, Sch04]), and if yes, return FALSE. Otherwise, still for $k > 1$, return the conjunction $\texttt{EMB}(X_1) \wedge \cdots \wedge \texttt{EMB}(X_k)$.

2. (Fill spherical holes) Now we have X connected. If it is an S^3, return TRUE. If there are components of ∂X that are S^2's, form X' by attaching a 3-ball to each spherical component of ∂X, and return $\texttt{EMB}(X')$.

3. (Connected sum) Form a decomposition $X = X_1 \# \cdots \# X_k$ of X into a connected sum[8] of prime manifolds[9] that are not 3-spheres.[10] If $k > 1$, i.e., X is not prime, return $\texttt{EMB}(X_1) \wedge \cdots \wedge \texttt{EMB}(X_k)$. If X is prime but not irreducible, i.e., contains an S^2 that does not bound a ball, then return FALSE.

4. (Boundary compression) Test if there is a compressing disc D for ∂X (i.e., $\partial D \subset \partial X$ does not bound a disk in ∂X).[11] If yes, cut X along D, obtaining a new manifold X'. Three cases may occur:

   (a) If X' has two components, X'_1 and X'_2, return $\texttt{EMB}(X'_1) \wedge \texttt{EMB}(X'_2)$. This case may occur, for example, for X a handlebody with two handles (a "thickened 8") when D separates the two handles.

   (b) If X' is connected and the two "scars" after cutting along D lie in the same component of $\partial X'$, return $\texttt{EMB}(X')$. This case may occur, e.g., for X a solid torus.

   (c) If neither of the previous two cases occur, then X' is connected but the scars lie in different components of $\partial X'$. Return FALSE. To get an example of X fitting this case, we can start with a thickened torus (i.e., torus times $[-1, 1]$) and connect the two boundary components by a 1-handle—which cannot be done in $\mathbb{R}^3$, but it does give a 3-manifold (with double torus boundary).

5. (Short meridian) Now X is irreducible and with incompressible boundary. Using [JR03, Thm. 5.20], retriangulate X with a 0-efficient triangulation. Then proceed as described at the beginning of Section 2: let $\gamma_1, \ldots, \gamma_n$ be a list of all closed essential normal curves in ∂X up to the length bound as in Theorem 1.4, for each i form $X'(\gamma_i)$ by attaching a 2-handle along γ_i, and return the disjunction $\texttt{EMB}(X'(\gamma_1)) \vee \cdots \vee \texttt{EMB}(X'(\gamma_n))$.

Lemma 3.1. *The above procedure always terminates and returns a correct answer, assuming the validity of Theorem 1.4.*

Proof. First we show that the algorithm always terminates. Let $C_1, \ldots, C_k$ be the components of ∂X numbered so that $g(C_1) \geq \cdots \geq g(C_k)$, where $g(.)$ stands for the genus, and let $\vec{g}_{\geq}(X)$ be the vector $(g(C_1), \ldots, g(C_k))$. We consider these vectors ordered lexicographically (if two vectors have a different length, we pad the shorter one with zeros on the right).

Let us think of the computation of the algorithm as a tree, with nodes corresponding to recursive calls. The branching degree is finite, so it suffices to check that every branch is finite.

It is easy to see that $\vec{g}_{\geq}(X)$ cannot increase by passing to a connected component or to a prime summand, and that it decreases strictly by a boundary compression and also by the short meridian step. Indeed, we observe that in the boundary compression step or the short meridian step, exactly one of the boundary components C_i is affected, and it is either split into two components C' and C'' of nonzero genus and with $g(C_i) = g(C') + g(C'')$, or it

[8] For two 3-manifolds X and Y, the *connected sum* $X \# Y$ is obtained by removing a small ball from the interior of X, another small ball from the interior of Y, and identifying the boundaries of these two balls.

[9] A *prime* 3-manifold is one that has no decomposition as a connected sum $X \# Y$ with neither X nor Y an S^3.

[10] The algorithm for prime decomposition goes back to Schubert [Sch49], for closed manifolds it is presented in detail in [JT95], and a version for manifolds with boundary is implicit in [JR03].

[11] The idea of an algorithm is due to Haken, and the algorithm is implicit in [JR03].

remains in one piece but the genus decreases by one. Since after steps 1–3 we have a connected irreducible manifold without spherical boundary components, for which the next step either finishes the computation or reduces $\vec{g}_{\geq}(X)$ strictly, every branch is finite as needed.

It remains to show that the returned answer is correct. For Step 2, we need that there is a unique way of filling a spherical hole; this is well known and can be inferred, for example, from the fact that there is only one orientation-preserving self-homeomorphism of S^2 up to isotopy [FM11, Sec. 2.2].

For Step 3, it is easily checked that a connected sum embeds iff the summands do. Moreover, every S^2 embedded in S^3 separates it, and hence if X contains a non-separating S^2, then it is not embeddable.

For Step 4, it is clear that if X is embeddable, then so is X'.

If, in case (4a), X'_1 and X'_2 are both embedded, then it is easy to construct an embedding of X: Denote D's scars by D_1 and D_2. Then a regular neighborhood of D_i is a ball B_i with boundary $S_i = \partial B_i$, and that meets both X'_i and $S^3 \setminus X'_i$ in balls. Think of each X'_i as embedded in its own copy of S^3, and take a connected sum of these two S^3's so that $S^3 = S^3 \# S^3 \supset X'_1 \#_{D_1=D_2} X'_2 = X$. Similarly, if X' is embedded in case (4b), then we can connect the scars by a thin handle in $S^3 \setminus X'$ and obtain an embedding of X.

In case (4c), let $C_1 \neq C_2$ be the components of $\partial X'$ containing the scars. Since the disk D does not separate X, we can choose a loop $\delta \subset$X meeting D in a single point, and such that δ also meets C_1 in a single point. But then, if X were embedded in S^3, C_1 would yield a nonseparating surface in S^3—a contradiction.

Finally, if one of the $X'(\gamma_i)$ is embeddable in Step 5, then so is X (since in X' we have 2-handle that was added to X, and we can just assign it to the complement of X), and if X is embeddable, then at least one of the $X'(\gamma_i)$ is by Theorem 1.4. □

4 Intersections of curves and surfaces

In this section we collect terminology, definitions and basic results concerning properly embedded curves in surfaces and properly embedded surfaces in 3-manifolds. In particular, for latter sections we need that any pair of properly embedded surfaces, each either essential, or, strongly irreducible and boundary strongly irreducible, can be isotoped to intersect essentially. There are few new results in this section. The reader is referred to Hempel [Hem76] and Jaco [Jac80] for more background.

We assume throughout that all curves and surfaces have been isotoped to have transverse intersection.

4.1 Curves

A *curve* is a properly embedded 1-dimensional manifold in a surface F, each component either a *loop*, which is closed, or an *arc*, which has two endpoints in ∂F.

A loop is *trivial* if it bounds a disk in F and an arc is *trivial* if it co-bounds a disk in F with some arc in ∂F. A curve is *essential* if no component is trivial.

Pairs of curves are assumed to intersect transversally. If α and β are a pair of curves, then their *geometric intersection number* $i(\alpha, \beta) = \min(|\alpha' \cap \beta'|)$ taken over all pairs of curves (α', β') where α' and β' are isotopic to α and β within F, respectively. (The isotopies are also allowed to move endpoints of arcs within the boundary.)

We say that α and β bound a *bigon* if there is a disk bounded by a pair of sub-arcs, one from each curve; see Figure 3 in Section 6 below. We say that they bound a *half-bigon* if there is a disk bounded by a pair of sub-arcs, one from each curve, along with an arc in ∂F. If α and β bound a bigon or half-bigon, then they can be isotoped to reduce their intersection.

We need this converse, a mild generalization of Farb and Margalit's *bigon criterion*:

Lemma 4.1 (Bigon criterion [FM11])**.** *A pair of curves α and β realize their geometric intersection number if and only if they do not bound a bigon or half-bigon.*

Proof. Farb and Margalit show that any pair of connected loops that intersect non-minimally form a bigon. They also note that this extends to disconnected curves consisting of loops.

If either curve has an arc component, then the doubled curves are properly embedded closed curves in the double[12] of the surface. If they intersect non-minimally in the original, they intersect non-minimally in the double and hence bound a bigon there. Thus, they bound a half-bigon in the original. □

4.2 Essential surfaces

We will assume that our surfaces are properly embedded in a 3-manifold X that is *irreducible*, i.e., every sphere embedded in X bounds a ball in X, and *boundary incompressible*, i.e., any curve in ∂X bounding a disk in X also bounds a disk in ∂X (is trivial).

Let F be a surface properly embedded in X. A *compressing disk* for a F is an embedded disk $D \subset X$ whose interior is disjoint from F and whose boundary is an essential loop in F. A *boundary compressing disk* is an embedded disk $D \subset X$ whose boundary, $\partial D = f \cup x$, is the union of $f = \partial D \cap F = D \cap F$, an essential arc properly embedded in F, and $x = \partial D \cap \partial X = D \cap F$, an arc properly embedded in ∂X. Here is an illustration:

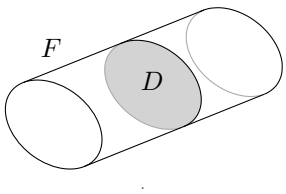

compressing disk for F

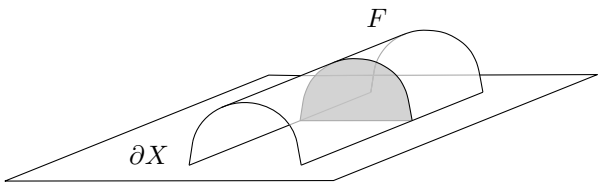

boundary compressing disk for F

A surface F is *compressible* if it has a compressing disk, *boundary compressible* if it has a boundary compressing disk, and *incompressible* and *boundary incompressible* if not, respectively. A surface is *essential* if it is incompressible, boundary incompressible, and not a sphere bounding a ball, or a disk co-bounding a ball with a disk in ∂X.

We establish some basic facts about surfaces in X.

Proposition 4.2. *The following statements hold for properly embedded surfaces in X, an irreducible, orientable 3-manifold with non-empty incompressible boundary:*

(i) *Every disk co-bounds a ball with a disk in ∂X.*

(ii) *Every connected surface with an inessential boundary curve is either compressible or a disk.*

[12] Meaning that we glue two copies of F by identifying their boundaries.

(iii) *The boundary curve of every compressible annulus is trivial in ∂X boundary.*

(iv) *Every boundary compressible annulus is boundary parallel (parallel to an annulus in ∂X).*

(v) *No surface is a projective plane.*

(vi) *Every Möbius band is essential.*

Proof. Because X has incompressible boundary, the boundary of a properly embedded disk bounds a disk in ∂X. The union of these disks is a sphere that, because X is irreducible, bounds a ball, yielding (i).

For (ii), suppose that some boundary curve of connected F bounds a disk in ∂X. Then, among disks in ∂X bounded by boundary curves of F, an innermost such disk can be pushed slightly into the interior of X while keeping its boundary in F. The boundary curve is either trivial in F, in which case F is a disk, or essential in F, in which case F is compressible.

Concerning (iii), let D be a compressing disk for an annulus A. Then ∂D separates A into two annuli A' and A''. So $D \cup A'$ and $D \cup A''$ are properly embedded disks, each with one boundary curve of A. Because ∂X is incompressible, the boundary curves are both trivial in ∂X.

As for (iv), let B be a boundary compressing disk for an annulus A. Then $\partial N(A \cup B)$, the boundary of a regular neighborhood of their union, has two components, an annulus isotopic to A and a disk. By (i), the disk co-bounds a ball with a disk in ∂X. But then the union of $N(A \cup B)$ with the ball is a solid torus, across which A is parallel to an annulus in ∂X.

Concerning (v), if P is a projective plane, then $\partial N(P)$, the boundary of its regular neighborhood, is a sphere which separates P from ∂X. Then X is reducible, for the sphere cannot bound a ball—no ball has interior boundary or contains an embedded projective plane.

Finally, in (vi), let M be a a Möbius band. Suppose first that M is compressible and let D be a compressing disk for M. Then ∂D cannot meet M in a core curve of M for this would imply that the core curve is orientation reversing in X. So ∂D is a 2-sided curve in M and separates it into an annulus and a narrower Möbius band M'. Then the union $M' \cup D$ is an embedded projective plane contradicting (v).

Suppose a Möbius band M is boundary compressible. Then $\partial N(M)$ is a boundary compressible annulus A. By (iv), A is boundary parallel, and co-bounds an solid torus with an annulus in the boundary. But the parallel region cannot contain the Möbius band M, and hence X is the union of two solid tori, $N(M)$ and the solid torus parallel region. □

We say that a a pair of surfaces, F and G, *intersect essentially* if each component of the curve $F \cap G$ is essential in both F and G (they are allowed to be disjoint). It is well known that essential surfaces can be arranged to intersect essentially:

Lemma 4.3. *Let F and G be properly embedded essential surfaces in an irreducible manifold with incompressible boundary. Then G can be isotoped so that they intersect essentially.*

Proof. Assume that we have isotoped G to minimize the number of curve components in $F \cap G$. We will show by contradiction that F and G intersect essentially.

We first note that if there is an intersection curve that is inessential in F, then there is an intersection curve that is inessential in G and vice-versa: If an intersection curve bounds a disk in F, choose one whose disk is innermost. Since G is incompressible, this disk is not a

compressing disk for G and it follows that its boundary, an intersection loop, is inessential in G. The same observation applies to inessential intersection arcs.

Then, assuming that some intersection loop is trivial, we can pass to one that is innermost on F, i.e., choose α to be an intersection loop that bounds a disk $D \subset F$ whose interior is disjoint from G. Since G is not compressible, α also bounds a disk $D' \subset G$. The union $D \cup D'$ is a sphere that, because X is irreducible, bounds a ball. And there is an isotopy of G that is restricted to a neighborhood of D', and that pushes D' across the ball and past D. This isotopy of G eliminates α and any other intersection curves in the interior of $F \cap D'$, and it does not introduce any new intersection curves since α was innermost.

Now assume some intersection arc is trivial in one of the surfaces, and as noted, we can let α denote such an arc that is outermost in F. That is, α cuts off a disk $D \subset F$ whose interior is disjoint from G and whose boundary meets ∂X in an arc. And $\alpha = D \cap G$ cuts off a, not necessarily outermost, disk $D' \subset G$ that also meets ∂X in an arc.

The union $D \cup D'$ is a disk with its boundary in ∂X that, because ∂X is incompressible, bounds a disk $D'' \subset \partial X$. Since X is irreducible, $D \cup D' \cup D''$ is a sphere bounding a ball. Moreover, there is an isotopy of G that pushes a neighborhood of D' past D and outside the ball. □

4.3 Almost meridional surfaces

Suppose that X is an irreducible manifold with incompressible boundary that is embedded in S^3. We recall that an essential curve $\mu \subset \partial X$ is a *meridian* if it bounds a disk in $S^3 \setminus X$. A properly embedded surface is *meridional* if each of its boundary curves is a meridian, and *almost meridional* if all but exactly one of its boundary curves is a meridian.

Let D be a boundary compressing disk for an orientable surface P. Then $\partial N(P \cup D)$ is a surface with at least two components. One component is isotopic to P; let P' be the union of the other components. Then P' is said to be the result of *boundary compressing* P along D.

Lemma 4.4. *Suppose that a manifold X is embedded in S^3. If P is a connected almost meridional planar surface properly embedded in X, then any surface P' obtained by boundary compressing P contains an almost meridional component.*

Proof. Let P' be obtained from P by boundary compressing along the disk D. What happens to ∂P? The disk D meets at most two boundary components of ∂P. Any component not met by D has two parallel copies in $\partial N(P \cup D)$, one for P and one for P', so those are unchanged. Let β be the one or two loops meet by the arc $x = D \cap \partial X$. Since D lies on one side of the the 2-sided planar surface P, when β is a single loop, x approaches it twice from the same side. It follows that $\partial N(x \cup \beta)$ is a *pair of pants*, i.e., an S^2 with three holes bounded by loops. One of these loops belongs to P and two to P', or vice-versa.

If any two of these three loops are meridians, then so is the third, since it bounds a disk, namely the union of the pants and the two disks pushed slightly into $S^3 \setminus \text{interior}(X)$.

We apply this "two meridians implies three meridians" principle to show that P' has an almost meridional component, regardless of how the boundary compressing disk meets the boundary components of P.

If the boundary compressing disk meets the non-meridional component twice, then the compression eliminates the non-meridional curve, and creates two new curves, each belonging to a separate component of P'. At least one of the new curves is not meridional, and hence its component is almost meridional.

If the boundary compressing disk meets a meridian and the non-meridian, then the compression does not separate P, and trades these curves for a new non-meridional curve. Thus P' is almost meridional.

If the boundary compressing disk meets two distinct meridians, then they are eliminated and a new one is created. The connected surface P' is almost meridional.

If the boundary compressing disk meets a single meridian twice, then P' has two components, each with one of the two new curves, either both meridional or both non-meridional. If both are meridional, then the component with the original non-meridian on its boundary is almost meridional. If both are non-meridional, then the component without the original non-meridian is almost meridional. One of the two components of P' is almost meridional. □

Lemma 4.5. *Suppose that X, an irreducible manifold with incompressible boundary, is embedded in S^3. If X contains an incompressible, almost meridional planar surface, then X contains an essential almost meridional planar surface.*

Proof. An incompressible almost meridional surface can be sequentially boundary compressed until it is incompressible and boundary incompressible. By the prior lemma, each surface in the sequence, hence the final one, has an almost meridional component. This final component is not a disk because X is boundary incompressible. Hence it is an essential almost meridional planar surface. □

4.4 Strongly irreducible surfaces

A two-sided surface properly embedded in X is *bi-compressible* if it has a *compressing pair* (D_+, D_-), a pair of disks, each a compressing or boundary compressing disk, one for each side of the surface. The pair is *simultaneous* if $\partial D_+ \cap \partial D_- = \emptyset$.

A surface is *weakly reducible* if it is simultaneously bi-compressible using compressing disks only. A *strongly irreducible* surface is one that is bi-compressible using compressing disks but not simultaneously so. A surface is *boundary weakly reducible* if it is simultaneously bi-compressible using any combination of compressing disks and boundary compressing disks. A surface is *boundary strongly irreducible* if it is bi-compressible, using any combination of compressing or boundary compressing disks, but not simultaneously so.

Some of our results assume that a surface is both strongly irreducible *and* boundary strongly irreducible. It may seem that the strongly irreducible hypothesis is vacuous. But this is not the case—it guarantees that the surface has at least one (non-boundary) compressing disk for each side.

Lemma 4.6 ([BDTS12], Lemma 3.8)**.** *In an irreducible manifold with incompressible boundary, the boundary of a strongly irreducible surface is essential in the boundary of the manifold.*

We state here a special case of Lemma 4.2 of [Bac13]. This

Lemma 4.7 (Lemma 4.2 of [Bac13])**.** *Let F be an essential surface and G a surface that is strongly irreducible and boundary strongly irreducible. Then G may be isotoped so that F and G intersect essentially.*

Let us remark that Bachman does not give a proof but claims it to be a direct generalization of [Bac09, Corol. 3.8]. He has also provided us with an unpublished manuscript with a proof.

5 Theory of normal curves and surfaces in a marked triangulation

In this section we introduce a mild generalization of the theory of normal curves and surfaces.

Definition 5.1. *A* marked triangulation *is a pair* $(\mathcal{T}, M)$ *consisting of a triangulation* $\mathcal{T}$ *of a* 2- *or* 3*-manifold along with a* marking $M \subset \mathcal{T}^1$*, a finite set of points along the edges of* $\mathcal{T}$*.*

If $M = \emptyset$, then $(\mathcal{T}, M)$ is a triangulation in the usual sense and we will usually omit M and refer directly to $\mathcal{T}$. Similarly, when $M = \emptyset$, we will describe objects as being normal rather than M-normal, and note that our definitions restrict to the standard ones.

An arc in a triangle is *M-normal* if its endpoints lie in distinct edges of the face and it misses M. A properly embedded curve α in the boundary surface is *M-normal* if it is the union of M-normal arcs. The length of α, $\ell(\alpha) = |\alpha \cap \mathcal{T}^1|$, is its number of intersections with the 1-skeleton.

There are several types of elementary surfaces contained in a tetrahedron Δ. An *M-normal disk* is a disk in Δ whose boundary is an M-normal curve of length 3 or 4 in $\partial\Delta$. We also consider two types of *M-exceptional pieces*: An octagon is a disk in Δ whose boundary is an M-normal curve of length 8 in $\partial\Delta$. A *tube* is an unknotted annulus in Δ whose boundary consists of two M-normal curves whose total length is at most 8.

An *M-normal surface* is a properly embedded surface that is the union of M-normal disks. An *almost M-normal surface* is a properly embedded surface that is the union of a single M-exceptional piece and a collection of M-normal disks.

The weight of an (almost) M-normal surface A is $\mathrm{wt}(A) = |A \cap \mathcal{T}^1|$, the number of intersections with the 1-skeleton. Its length is the length of its boundary: $\ell(A) := \ell(\partial A)$.

An *M-normal isotopy* is a normal isotopy that does not pass through any point in M. An *M-type* is the equivalence class of an M-normal arc in a face, or, an M-normal disc or M-exceptional piece in a tetrahedron. Two types are *M-compatible* if they have disjoint representatives. A pair of curves or surfaces are *M-compatible* if each pair of types they possess are M-compatible. That is, a pair of curves is M compatible if, for each face of the triangulation, their arcs in that face are pairwise disjoint after an M-normal isotopy. An analogous statement holds for surfaces.

We note that M-compatibility is a local condition; in general it may not be possibly to make M-compatible curves or surfaces globally disjoint by an M-normal isotopy.

The *M-normal vector* or *M-normal coordinates* of an M-normal curve, surface, or almost M-normal surface A is a uniquely determined vector $\vec{v}_M(A)$, indexed over the set of normal types and with each entry recording the number of M-normal objects of the index type.

If A, B, C are M-normal surfaces such that $\vec{v}_M(C) = \vec{v}_M(A) + \vec{v}_M(B)$, then C is an *M-normal sum* of A and B, and we write $C = A + B$. The same definition applies if A is M-normal and B and C almost M-normal, or if A, B, C are M-normal curves.

We note that not every two M-normal surfaces, for example, can be normally added—this is possible exactly if they are M-compatible.

If A and B are M-compatible, then one can construct an M-normal sum as follows. In each face or tetrahedron Δ, the M-normal pieces $A \cap \Delta$ and $B \cap \Delta$ can be M-normally isotoped to be disjoint, and then attached across each facet of Δ to the pieces in an adjacent face/tetrahedron. This produces a properly embedded M-normal curve, M-normal surface, or almost M-normal surface, respectively, which is the M-normal sum.

However, in our considerations, we will mostly use a different geometric construction of an M-normal sum, where we assume that the curves or surfaces in question intersect minimally, in a suitable sense, but then we do not isotope them to be disjoint as above, but rather they stay in place and we deal with their intersections as well; see Section 5.1 below.

It is well known that Euler characteristic, weight and length are all additive with respect to normal sum, and this works without change for the M-normal case. If A and B are compatible (almost) M-normal curves or surfaces then the following hold:

1. $\chi(A + B) = \chi(A) + \chi(B)$
2. $\mathrm{wt}(A + B) = \mathrm{wt}(A) + \mathrm{wt}(B)$
3. $\ell(A + B) = \ell(A) + \ell(B)$.

An (almost) M-normal curve or surface is *fundamental* if it cannot be expressed as the sum of other (almost) M-normal curves or surfaces. Every (almost) M-normal curve/surface is a non-negative integer combination of fundamentals.

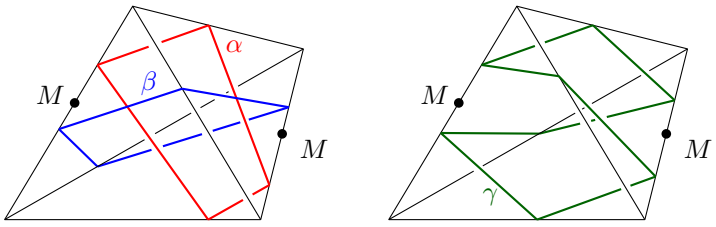

$\vec{v}(\gamma) = \vec{v}(\alpha) + \vec{v}(\beta)$ $\qquad$ $\vec{v}_M(\gamma) \neq \vec{v}_M(\alpha) + \vec{v}_M(\beta)$

Figure 1: A length 8 curve that is not the sum of two length 4 curves in the marked triangulation.

Here M-normal curve theory differs from standard normal curve theory. While all normal curves are compatible, M-normal curves have distinct compatibility classes, and this increases the number of fundamentals. In Figure 1, we see the boundary of a tetrahedron with two marked points, one on each of a pair of opposite edges. Let γ be the length 8 M-normal curve that meets each of the sub-edges once. As a normal curve γ is not fundamental—it is the sum of the two distinct length 4 curves α and β. But in the marked triangulation these curves are incompatible and γ is fundamental.

If α is an (almost) M-normal curve or surface, then $\vec{v}_M(\alpha)$ is a solution to a set of *matching equations*: For a triangulated surface, this set consists of one equation for each sub-edge in the interior of the surface. It sets equal the sum of those coordinates meeting the sub-edge from one side to the sum of those meeting it on the other side. In a triangulated 3-manifold,

the set of matching equations consists of one equation for each M-normal arc type contained in an interior face. This equation sets equal the sum of the coordinates for elementary types using the arc type on one side to the sum of those using it on the other side.

We say that a vector $\vec{v}_M$ of the correct dimension is *M-admissible* if all its coordinates are non-negative, it satisfies the matching equations, and is *self-compatible*, i.e., it does not possess non-zero coordinates for any pair of non-M-compatible types. If $\vec{v}_M$ is an M-admissible vector, there is an (almost) M-normal curve/surface α for which $\vec{v}_M = \vec{v}_M(\alpha)$.

The following proposition is a straightforward generalization of a well known fact from normal surface theory to M-normal surfaces; see [HLP99] for a nice exposition.

Proposition 5.2. *Given a 3-manifold with a marked triangulation $\mathcal{T}_M$, the set $\mathcal{F}$ of fundamental M-normal surfaces is computable, and both $|\mathcal{F}|$ and the maximum weight* $\mathrm{wt}(F)$ *of an $F \in \mathcal{F}$ are bounded by a computable function of t and m.*

The bound on $|\mathcal{F}|$ and $\mathrm{wt}(F)$ has the form $\exp(p(t,m))$, where $p(t,m)$ is a suitable polynomial.

Proof. It is well known that without the marking, there are $7t$ normal disk types, $3t$ exceptional octagons and $25t$ exceptional tubed pairs of disks. Moreover, the presence of a tubed pair of disks may split one type of normal disks into two, but certainly we have no more than $42t$ types in total.

The points of M divide each edge into at most $m+1$ subarcs. In order to specify an M-normal type of a triangle, for example, we need to specify the subarc containing each of the three vertices, which leads to the bound $(m+1)^3$. The worst bound is obtained for tubes and octagons, with $(m+1)^8$, so a rough bound for the total number of M-types is $42t(m+1)^8$.

A similar way of counting applies to the number of *matching equations*, which represent compatibility of the coordinates of the M-normal vector across the pieces of the edges of $\mathcal{T}$ delimited by the points of M. Indeed, the matching equations correspond to M-arc types. There are at most $4t$ interior faces, each with 3 underlying normal arc types. A given M-arc type is thus determined by this normal type and by the sub-arcs it meets, and so there are at most $12t(m+1)^2$ matching equations.

Then, reasoning as in [HLP99, Sec. 6], using a Hilbert basis of the appropriate integral cone, we obtain the bounds of the claimed form. □

5.1 Snug pairs of curves and surfaces, Haken sums, and normal sums

The normal sum $F+G$ of a pair of (almost) M-normal curves or surfaces F and G has been defined, if they are M-compatible, to be an M-normal surface whose M-normal vector is the sum of the M-normal vectors of F and G, $\vec{v}_M(F+G) = \vec{v}_M(F) + \vec{v}_M(G)$.

It is desirable to show that qualities of the sum, such as essentiality or minimality, also apply to the summands. Here we describe a well known geometric interpretation of the sum that makes this possible; also see, for example, [JO84, JT95]. We also present some related material.

Snug pairs. We begin with a definition of a "placement with no unnecessary intersections" for a pair of curves or surfaces.

Definition 5.3. *A pair (F,G) of properly embedded curves or surfaces is* snug *if it is transverse and the number of components of the intersection $F \cap G$ is minimized over pairs (F',G'),*

where F' and G' are isotopic to F and G, respectively. The pair (F, G) is locally snug *if $F \cap G$ is disjoint from the 1-skeleton $\mathcal{T}^1$, and, they are snug in the interior of each simplex of the triangulation (here we only allow isotopies moving each intersection of F or G with a face only within that face).*

If F and G are locally snug M-normal surfaces then it follows that:

1. each pair of M-normal arcs, one from F and one from G, meets in 0 or 1 points;
2. each pair of M-normal disks, one from F and one from G, meets in 2 or fewer arcs, and the union of the arcs has at most one endpoint in any face;
3. no loop of $F \cap G$ lies inside a tetrahedron.

Any pair of compatible M-normal curves or surfaces can be made locally snug by M-normal isotopies that first make their intersections with edges disjoint and then "straighten" them so that: normal arcs are straight, normal triangles are flat, and normal quads are the union of two flat triangles. We do not define locally snug when F is an almost normal surface and G is a normal surface, for in that case we require only the definition of the normal sum $F + G$ and not its geometric interpretation.

Haken sum and normal sum of curves. Now, for a while, we deal only with curves, and we develop a geometric interpretation of their normal sum. Here we consider only unmarked triangulations, i.e., $M = \emptyset$.

Let D be a regular neighborhood of an intersection point x of a pair of transverse curves α and β. We can remove the intersection by deleting the arcs in the interior of the disk and then attaching α to β along a pair of antipodal sub-arcs of ∂D. Thus, we replace the "$\times$" in $\alpha \cup \beta$ with either "$)($" or "$\asymp$". This is called an *exchange* or a *switch* at x. A curve is said to be a *Haken sum* $\alpha \mathbin{\widetilde{+}} \beta$ of α and β if it is obtained by an exchange at each of their intersection points. Of course, $\alpha \mathbin{\widetilde{+}} \beta$ is dependent on the direction of the switches and is therefore not well determined.

If, however, α and β are locally snug normal curves, then each intersection point is of the form $x = \alpha' \cap \beta'$ where α' and β' are normal arcs in some face. Then α' and β' meet at least one common edge e of the face. The *regular exchange* is the exchange that does not produce an *abnormal arc*, a non-normal arc with both endpoints attached to e; see Fig. 2 top.

As we will see, the normal sum of $\alpha + \beta$ of locally snug curves can be obtained by doing all the regular exchanges.

Lemma 5.4. *Let α and β be locally snug normal curves. Then the Haken sum $\alpha \mathbin{\widetilde{+}} \beta$ obtained by making all the regular exchanges is the normal sum $\alpha + \beta$; i.e., $\vec{v}(\alpha + \beta) = \vec{v}(\alpha) + \vec{v}(\beta)$.*

Proof. We show that the result holds in each face of the triangulation. In an abuse of notation, let α and β be restriction of the curves to a particular face. For contradiction suppose that they are a counterexample that minimizes $|\alpha \cap \beta|$. Then α and β are not disjoint, for in that case, the union is normal and normal vectors add.

Since they intersect in a face, we can identify an outermost half-bigon bounded by a sub-arcs of α and β and an edge of the face; see Figure 3. The regular exchange trades these sub-arcs and results in a pair of normal curves, α' normally isotopic to α and β' normally isotopic to β, that are locally snug but with fewer intersections. By assumption, these α' and β' satisfy the conclusion, hence so do α and β. □

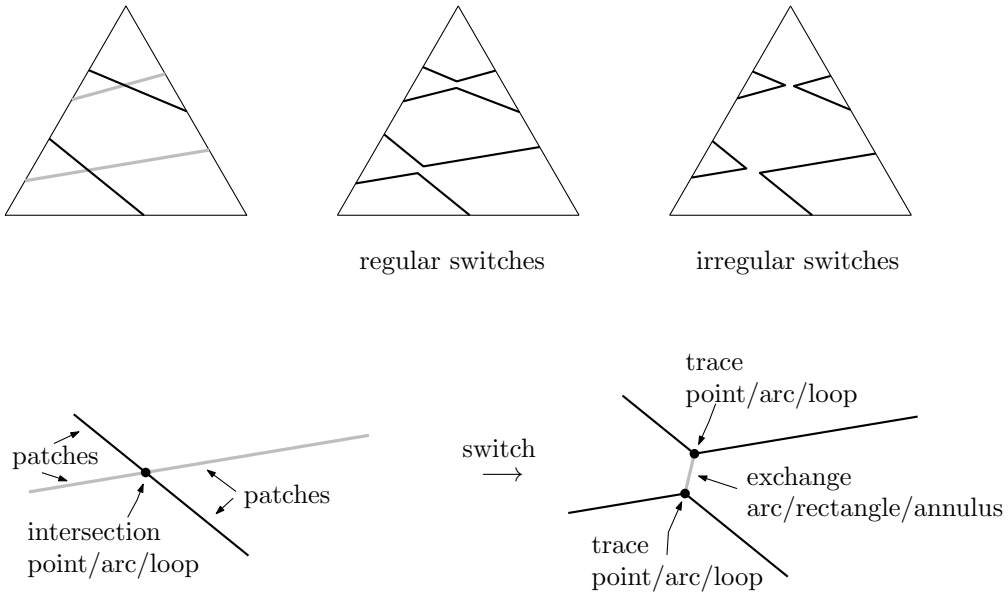

Figure 2: Switches for curves and surfaces

Lemma 5.5. *Let $\alpha \mathbin{\widetilde{+}} \beta$ be a Haken sum of locally snug properly embedded curves. Then $\alpha \mathbin{\widetilde{+}} \beta$ is normal if and only if α and β are normal and all switches are regular, i.e., $\alpha \mathbin{\widetilde{+}} \beta = \alpha + \beta$. In addition, if α and β are normal and $\alpha \mathbin{\widetilde{+}} \beta$ contains at least one irregular switch, then $\alpha \mathbin{\widetilde{+}} \beta$ contains an abnormal arc.*

Proof. ($\Leftarrow$) This is by Lemma 5.4. ($\Rightarrow$) We show that if either α or β is not normal then neither is $\alpha \mathbin{\widetilde{+}} \beta$ for any Haken sum of the curves.

Suppose then that $\alpha' \subset \alpha$ is an outermost non-normal arc, one that co-bounds a disk with a sub-arc of an edge $e' \subset e$. If β meets the disk, it meets it in a collection of n arcs, each with one endpoint in α' and one endpoint in e', because α' is outermost and α and β are snug. Let D be a regular neighborhood of the disk. Then, regardless of the switches, $\alpha \mathbin{\widetilde{+}} \beta$ meets D in a collection of $n+1$ arcs that have $n+2$ endpoints along the edge and n endpoints not on the edge. It follows that at least one arc meets the edge in 2 points and is not normal. A symmetric argument applies if the outermost non-normal arc belongs to β. Nor can either α or β possess a loop in a face. Local snugness implies that any loop is disjoint from the other curve and survives any Haken sum.

We now know that α, β are normal. To conclude the proof, it is sufficient to show that $\alpha \mathbin{\widetilde{+}} \beta$ contains an abnormal arc if at least one switch is irregular (this contradicts the normality of $\alpha + \beta$, and thus proves the last claim of the lemma). In an abuse of notation, let α and β refer to the collection of normal arcs in a particular face.

We perform the specified switches in order according to the following scheme: If $\alpha \cap \beta \neq \emptyset$, then α and β form an outermost half bigon B with an edge as in Figure 3. The regular switch produces collections α' and β' that are normally isotopic to α and β, but with one fewer intersections. An irregular switch produces a disjoint abnormal arc that survives any and all additional exchanges. If each exchange is regular, we can continue and the process produces a disjoint union $\alpha \sqcup \beta$. If any exchange is not regular, the resulting curve contains an abnormal arc. □

Figure 3: Exchanges at the corners of a bigon and a half-bigon.

Normal sign. When α and β lie in an oriented surface, for example the boundary of an oriented manifold, we can define the *normal sign* of each point of $\alpha \cap \beta$. Viewing α as horizontal and β as vertical, the regular exchange at the point connects a pair of quadrants. The point has *positive sign* if the exchange connects the southwest quadrant to the northeast quadrant, and it has *negative sign* if it connects the northwest to the southeast; see Figure 4. This is equivalent to the definition given in [BDTS12]. The definition depends on the ordering of the pair of curves and on an orientation on the surface: reversing the order or the orientation reverses every sign.

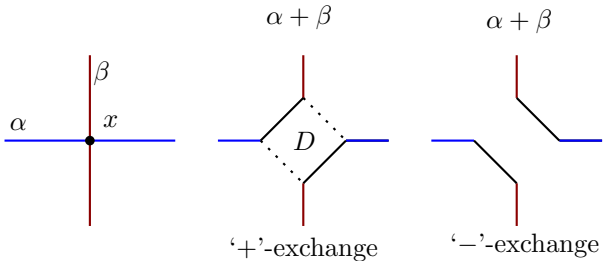

Figure 4: Exchanges with positive and negative sign.

Normal sum of surfaces. Similar to the case of curves above, one can also construct the normal sum $F + G$ of normal surfaces geometrically, using suitable switches. We assume that F and G are locally snug.

We construct $F + G$ by specifying its intersection with the 1-, 2-, and 3-skeleta of the triangulation, respectively. First, we let the intersection of $F + G$ with the 1-skeleton to be the union of the intersection points from F and those from G.

Second, in each face we perform regular switches on all intersecting pairs of arcs (f, g), where f comes from F and g from G.

Finally, we construct the normal sum $F + G$ in the interior of each tetrahedron T. As discussed earlier, each normal disk is either a flat triangle or a quadrilateral made of 2 flat triangles. It follows that every intersection between normal disks from compatible surfaces is either 1 or 2 arcs, not necessarily straight. Compatibility ensures that the regular switches prescribed at the endpoints of each arc are consistent with each other and can be extended across the entire arc of intersection. The normal sum F and G is the result of performing such regular switches along every such arc of intersection.

Note that any intersection arc between normal disks can be extended from a tetrahedron through a face to a neighboring tetrahedron. In its entirety this *intersection curve* between F and G is either a loop, or an arc with both endpoints in ∂X. Compatibility ensures that the regular switches in each face and through the interior of each tetrahedron agree. Thus we can regard the switch as a regular switch along the entire intersection curve.

Exchange arcs and surfaces, trace curves. Here we introduce some additional terminology. First, we consider a regular switch of two curves. Inside the neighborhood where the regular switch was performed, we identify an *exchange arc* that connects the points of the newly formed arcs corresponding to the former intersection points; see Figure 2.

Next, we consider two locally snug normal surfaces F and G. A *patch* is a component of $F \cup G \setminus F \cap G$. Regular switches reconnect the patches, and *trace curves* are the seams between patches after performing regular switches along all intersection curves; see Figures 2 and 5.

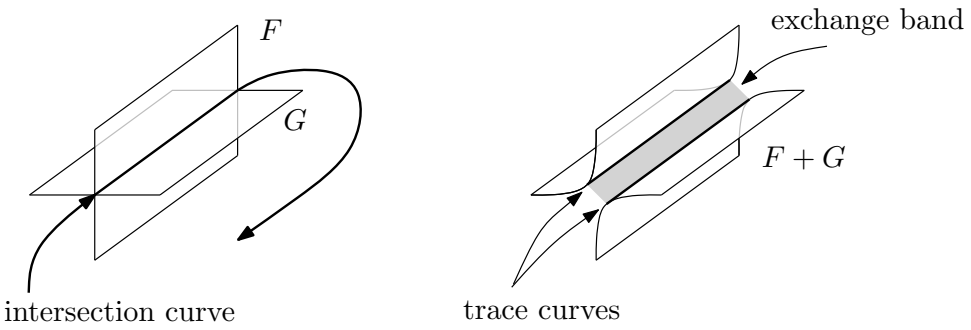

Figure 5: The exchange between surfaces along an intersection curve.

If the intersection curve α is an arc, then, after performing a regular switch, we can identify an *exchange rectangle*, a rectangle whose top and bottom, say, are bounded by trace arcs and whose left and right sides are *exchange arcs* lying in ∂X.

If α is a loop, then our assumption that X is orientable means that a regular neighborhood of α is a solid torus, not a solid Klein bottle. Again, since X is orientable, α is either orientation preserving in both F and G, or, orientation reversing in both F and G. In the former case, there is an *exchange annulus*, a zero-weight annulus bounded by the trace curves and with core α. In the latter case, there is a single trace curve which bounds an *exchange Möbius band* (we will be able to exclude this case in our proofs, though).

As observed in [Hat82] and, in the context of normal surfaces, in [JS03], every intersection arc between surfaces connects intersection points of the boundary curves that have opposite normal sign:

Lemma 5.6. *Let $A = B + C$ be a normal sum of surfaces in an orientable manifold X with an induced orientation on ∂X. Then every arc in $B \cap C$ joins a pair of points in $\partial B \cap \partial C$ with opposite normal sign.*

6 Complexity and tight curves

In this section, we consider properly properly embedded curves in a triangulated surface. We assume that they are transverse to the 1-skeleton but, a priori, they are not assumed to be normal.

Fix, once and for all, an ordering of all normal arc types of the triangulated surface. For this purpose we *do not* take into account any marking present. As in the previous section, a normal curve α determines a vector $\vec{v}(\alpha)$ which records the number of normal arcs of the indexed type. Order these normal vectors lexicographically.

Recall that the length of a properly embedded curve α is the number of intersections with the 1-skeleton, $\ell(\alpha) = |\alpha \cap \mathcal{T}^1|$. We say that a curve is *least length* if it minimizes length over all curves to which it is isotopic.

Lemma 6.1. *A least length essential curve is normal.*

Proof. A loop in face demonstrates that the curve is not essential and any abnormal arc is either inessential or yields an isotopy reducing the length. □

If α is a normal curve, then we define its *complexity* to be the pair consisting of its length and its normal vector,

$$\mathrm{cpx}(\alpha) := (\ell(\alpha), \vec{v}(\alpha)).$$

We reiterate that we do not take into account any marking M in the definition of complexity. If α is not normal, we define its complexity to be $\mathrm{cpx}(\alpha) = (\ell(\alpha), \vec{0})$. Complexities will also be ordered lexicographically.

Definition 6.2. *A curve α is* tight *if it minimizes the complexity* $\mathrm{cpx}(\alpha)$ *over all curves to which it is isotopic.*

The interior of a connected inessential curve can be made disjoint from the 1-skeleton, so a tight inessential loop has $\mathrm{cpx} = (0, \vec{0})$ and a tight inessential arc has $\mathrm{cpx} = (2, \vec{0})$.

Lemma 6.3. *A tight essential curve is normal and unique up to normal isotopy.*

Proof. Indeed, the complexities of two normal curves are equal if and only if their normal vectors are identical. □

Lemma 6.4. *Let $\alpha \mathbin{\widetilde{+}} \beta$ be a Haken sum of locally snug properly embedded curves. Then $\mathrm{cpx}(\alpha \mathbin{\widetilde{+}} \beta) \leq \mathrm{cpx}(\alpha) + \mathrm{cpx}(\beta)$ with equality holding if and only if α, β and $\alpha \mathbin{\widetilde{+}} \beta$ are all normal, or, all not normal.*

Proof. The curve $\alpha \mathbin{\widetilde{+}} \beta$ is constructed by performing an exchange at every intersection point of $\alpha \cup \beta$. This is done away from the 1-skeleton, so we have $\ell(\alpha \mathbin{\widetilde{+}} \beta) = \ell(\alpha) + \ell(\beta)$. Thus any difference in complexity is determined solely by the normal vectors of the curves. If $\alpha \mathbin{\widetilde{+}} \beta$ is normal, then by the previous two lemmas $\vec{v}(\alpha + \beta) = \vec{v}(\alpha) + \vec{v}(\beta)$ and equality holds. If $\alpha \mathbin{\widetilde{+}} \beta$ is not normal, then its normal vector is $\vec{0}$. Then complexity is additive when both α or β are not normal, and sub-additive otherwise. □

If a tight curve is written as a sum, then the exchange arcs are essential in the complement of the curve:

Lemma 6.5. *Suppose that a tight normal curve is written as a sum $\alpha + \beta$ of two normal curves. Then no exchange arc co-bounds a disk with a sub-arc of the curve.*

Proof. Perform an irregular exchange only at the intersection corresponding to this exchange arc. The new curve is a Haken sum $\alpha \mathbin{\widetilde{+}} \beta$ with one component a trivial loop, the rest isotopic to $\alpha + \beta$, and the same total length. It follows that the trivial loop has zero length otherwise $\alpha + \beta$ would not be tight. Therefore $\alpha + \beta$ and the second component of $\alpha \mathbin{\widetilde{+}} \beta$ are normally isotopic. However, by Lemma 5.5 there is an abnormal arc in the second component of $\alpha \mathbin{\widetilde{+}} \beta$, a contradiction. □

Lemma 6.6. *Let B be a bigon or half-bigon bounded by a pair of locally snug normal curves α and β; see Figure 3. Let α' and β' be the pair of isotopic curves obtained by corner exchange(s) that trade the sides of B. Then one of the following holds:*

(1) $\mathrm{cpx}(\alpha') = \mathrm{cpx}(\alpha), \mathrm{cpx}(\beta) = \mathrm{cpx}(\beta')$*, thus α' and β' are normally isotopic to α and β, respectively, and $|\alpha' \cap \beta'| < |\alpha \cap \beta|$;*

(2) $\mathrm{cpx}(\alpha') < \mathrm{cpx}(\alpha)$*;*

(3) $\mathrm{cpx}(\beta') < \mathrm{cpx}(\beta)$.

Proof. Let α, α', β and β' be as indicated in Figure 3. Note that the exchange doesn't add or remove intersections with the 1-skeleton, and so the total length is unchanged. If the traded arcs differ in length then one curve increases and the other decreases in length, hence in complexity. In this case, either (2) or (3) holds. So we continue assuming $\ell(\alpha) = \ell(\alpha')$ and $\ell(\beta) = \ell(\beta')$.

If any exchange is irregular, then one of the curves, say α', is not normal. Then its complexity $\mathrm{cpx}(\alpha') = (\ell(\alpha'), \vec{0}) = (\ell(\alpha), \vec{0}) < (\ell(\alpha), \vec{v}(\alpha)) = \mathrm{cpx}(\alpha)$ has decreased, yielding conclusion (2). Conclusion (3) results when β' is not normal.

We are left in the case that the exchange trades length fairly and α' and β' are both normal. Because length and normal vectors are both additive with respect to normal addition, we have $\mathrm{cpx}(\alpha)+\mathrm{cpx}(\beta) = \mathrm{cpx}(\alpha+\beta) = \mathrm{cpx}(\alpha')+\mathrm{cpx}(\beta')$. If $\mathrm{cpx}(\alpha') = \mathrm{cpx}(\alpha)$, then $\mathrm{cpx}(\beta') = \mathrm{cpx}(\beta)$ and by Lemma 6.3 the trade yields normally isotopic curves, conclusion (1). Otherwise, either (2) or (3) holds. □

Lemma 6.7. *Let α be a tight essential curve and $\mathcal{C}$ set of pairwise snug, tight essential curves. Then, after a normal isotopy of α, $\{\alpha\} \cup \mathcal{C}$ is pairwise snug.*

Proof. Normally isotope α to minimize the total of all intersections with $\mathcal{C}$. By way of contradiction, suppose some pair is not snug, that there is $\beta \in \mathcal{C}$ for which $|\alpha \cap \beta| > i(\alpha, \beta)$. Among all such β take one that, together with α, determines an innermost bigon; then any other curves from $\mathcal{C}$ meeting that bigon meet it in arcs that run straight across.

Apply Lemma 6.6. Since all curves are tight, we must have the first conclusion. But, trading across the bigon reduces intersections between α and β without raising intersections of any other pair—a contradiction. □

Lemma 6.8. *Suppose that a tight essential normal curve is a normal sum $\alpha + \beta$. Then α and β are tight, essential, and after a normal isotopy, snug.*

Proof. Normally isotope α and/or β to minimize $|\alpha \cap \beta|$. This does not change their sum.

First we show that the pair is snug: If not, then some pair of sub-arcs of α and β bound a bigon or half-bigon B. Apply Lemma 6.6. The first conclusion does not hold, so without loss of generality assume that $\mathrm{cpx}(\alpha') < \mathrm{cpx}(\alpha)$. Isotope α' back slightly so that α' and β still overlap and form a very thin bigon. Since $\alpha \cup \beta$ and $\alpha' \cup \beta$ are isotopic as graphs, $\alpha + \beta$ is isotopic to some Haken sum $\alpha' \mathbin{\widetilde{+}} \beta$. But by Lemma 6.4, $\mathrm{cpx}(\alpha' \mathbin{\widetilde{+}} \beta) \leq \mathrm{cpx}(\alpha') + \mathrm{cpx}(\beta) < \mathrm{cpx}(\alpha) + \mathrm{cpx}(\beta) = \mathrm{cpx}(\alpha + \beta)$, a contradiction.

It follows that α and β are both essential. If either possesses a component that bounds a disk, then the fact that α and β are snug implies that this component misses the other curve, survives normal addition, and $\alpha + \beta$ contains an inessential component, a contradiction.

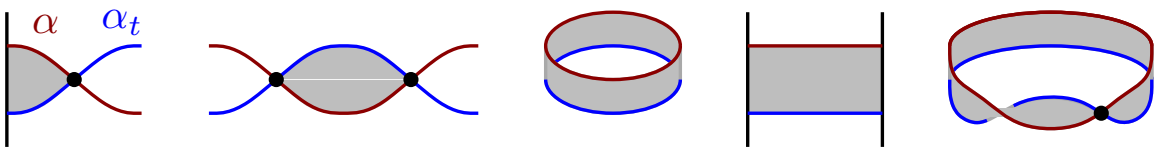

Figure 6: Possible product regions derived from isotopy of α and α_t.

It remains to show that each summand is tight. Without loss of generality, suppose that α is not tight, that there is a tight curve α_t with lower complexity, $\mathrm{cpx}(\alpha_t) < \mathrm{cpx}(\alpha)$, that is isotopic to α but not normally so. Isotope α_t to intersect $\alpha \cup \beta$ minimally.

Then any innermost (half-) bigon in the complement of $\alpha_t \cup \alpha \cup \beta$ is bounded by α and α_t, since it cannot be bounded by α and β, which are snug. And because any patch of β is a sub-arc of $\alpha + \beta$, any innermost (half-) bigon bounded by β and α_t is also a (half-) bigon bounded by the tight curves $\alpha + \beta$ and α_t which, using Lemma 6.6 again, can be eliminated by a normal isotopy of α_t. This contradicts the minimality of the intersection between α_t and $\alpha \cup \beta$.

Then, sub-curves of α and α_t co-bound a product region R as in Figure 6.

If they are not snug, R is a bigon or half-bigon. If they are snug, R is a rectangle when α is an arc, an annulus when α is a two-sided loop, and a bigon with corners identified when α is a one-sided loop.

In all of these cases, as observed above, no arc of β forms a (half-) bigon inside R, and must therefore run across R and have an endpoint in both α and α_t.

In the non-snug case, let α' be the curve of less complexity obtained by routing α along α_t when it meets the bigon or half-bigon. In the snug case, let $\alpha' = \alpha_t$. In either case, $\mathrm{cpx}(\alpha') < \mathrm{cpx}(\alpha)$. Moreover, the complex $\alpha' \cup \beta$ is isotopic to $\alpha \cup \beta$ and because they are isotopic, there are exchanges, not necessarily regular, so that the Haken sum $\alpha' \mathbin{\widetilde{+}} \beta$ is a curve isotopic to $\alpha + \beta$. But by Lemma 6.4, $\mathrm{cpx}(\alpha' \mathbin{\widetilde{+}} \beta) \leq \mathrm{cpx}(\alpha') + \mathrm{cpx}(\beta) < \mathrm{cpx}(\alpha) + \mathrm{cpx}(\beta) = \mathrm{cpx}(\alpha + \beta)$. This contradicts the fact that $\alpha + \beta$ is tight. □

Rails and fences. Now we again consider a triangulation with a marking M, and auxiliary curves in it that, unlike M-normal curves, go through the points of M.

A *rail* is a normal arc with its endpoints in M, and a *fence* is a normal curve that is the union of rails.

We note that if a face contains an M-normal arc α and a rail μ that are locally snug, then $|\alpha \cap \mu|$ is either 0 or 1, depending only on the endpoints of μ and the M-normal type of α.

The following lemma can also be considered obvious:

Lemma 6.9. *Intersection number with fences is additive with respect to normal addition of M-normal curves: If μ is an fence and α and β are M-compatible, M-normal curves, then $|(\alpha + \beta) \cap \mu| = |\alpha \cap \mu| + |\beta \cap \mu|$.*

Proposition 6.10. *Let μ be a fence that is a tight essential curve (w.r.t. the unmarked triangulation). Suppose that a sum $\alpha + \beta$ of M-normal curves is tight, essential and snug with μ. Then*

(1) *α and β are both snug with respect to μ;*

(2) *$i(\alpha + \beta, \mu) = i(\alpha, \mu) + i(\beta, \mu)$ where $i(.,.)$ is the geometric intersection number;*

(3) *if β is two-sided, connected and normally isotopic to μ then, after a normal isotopy, every point of $\alpha \cap \beta$ has the same normal sign.*

Proof. Among counterexamples to conclusion (1) of the proposition, choose one that minimizes $|\alpha \cap \beta|$. Then α, β and μ are pairwise locally snug, and we will show that they are in fact pairwise snug. Suppose not and let B be an innermost (half-) bigon bounded by some pair of the curves.

If B is bounded by μ and either of the other curves, say α, then every sub-arc of β in B crosses B and meets both α and μ. Let α' be the result of rerouting α around B as in Lemma 6.6. Then $\alpha + \beta$ is isotopic to some Haken sum $\alpha' \mathbin{\widetilde{+}} \beta$ that has fewer intersections with μ. This contradicts our assumption that $\alpha + \beta$ and μ are snug.

If B is bounded by α and β, then every sub-arc of μ in B crosses B and meets both α and β. Let α' and β' be the curves given by Lemma 6.6. Because $\alpha + \beta$ is tight, α and β are tight and normally, but not necessarily M-normally, isotopic to α' and β', respectively. Because the normal sum $\alpha + \beta$ was defined, α and β are locally snug. The isotopy doesn't create intersections, and so α' and β' are also locally snug. Then $\alpha + \beta = \alpha' \mathbin{\widetilde{+}} \beta'$ for some generalized Haken sum of α' and β'. By Lemma 5.5 that sum is a normal sum, $\alpha+\beta = \alpha'+\beta'$. Note that α' and β' are snug with μ if and only if α and β are, as the move did not change the number of times they meet μ. Since $|\alpha \cap \beta| > |\alpha' \cap \beta'|$ we obtain a contradiction and establish conclusion (1).

Since α, β and their sum are all snug with respect to μ and intersections with respect to μ are additive, we have additivity of geometric intersection number, conclusion (2).

We now prove the final statement of the proposition. Assume that β is normally isotopic to μ. Then $i(\beta, \mu) = 0$ since β is two-sided. By (2) and the fact that $\alpha + \beta$ and α are both snug with μ we have: $|(\alpha + \beta) \cap \mu| = i(\alpha + \beta, \mu) = i(\alpha, \mu) = |\alpha \cap \mu|$.

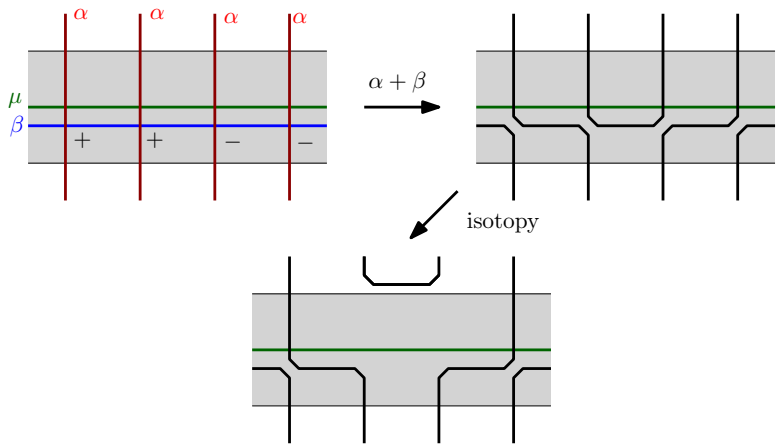

Figure 7: Mixed normal signs imply that $i(\alpha + \beta, \beta) < i(\alpha, \beta)$.

And since β is normally isotopic to μ, there is a normal, not necessarily M-normal, isotopy taking $\alpha + \beta$ to $\alpha + \mu$. Now α cuts across a thin regular neighborhood of μ in a collection of arcs that span (cut across) the neighborhood. Together they cut this neighborhood into rectangles; see Figure 7.

Each regular exchange connects a pair of rectangles at a corner of each. In fact, every rectangle that meets $\alpha \cap \mu$ twice must be attached to another rectangle at one of its corners.

Otherwise, an arc of $\alpha + \mu$ bounds the unattached rectangle, showing that the arc it is trivial in the neighborhood of μ and can be isotoped out of it. This would imply $i(\alpha+\beta, \mu) < |\alpha \cap \mu|$, contradicting the equality shown earlier. The fact that each rectangle is attached at exactly one corner implies that as we follow μ every intersection with α must have the same normal sign. Since β is normally isotopic to μ we have our desired conclusion (3). □

7 Normal summands of incompressible annuli

We would like to apply two well known results from normal surface theory: (1) an essential surface is isotopic to a normal surface, and (2) every summand of a least weight essential normal surface is also least weight and essential (Theorem 6.5 of Jaco and Tollefson [JT95]). But, as will be seen shortly, our notion of surface complexity prioritizes the reduction of boundary complexity over the reduction of total surface weight. Thus the results (1) and (2) cannot be applied as stated.

Proposition 7.1 recovers the first result using our notion of complexity. Proposition 7.2 gives a weaker version of the second for incompressible annuli. While we expect the full version to hold with our notion of complexity, we prove a restricted version both to simplify the proof and to incorporate boundary parallel annuli which are non-essential. Our proof follows the strategy of [JO84] and [JT95].

The *complexity* $\mathrm{cpx}(F)$ of a properly embedded surface F is the triple

$$\mathrm{cpx}(F) = (\mathrm{cpx}(\partial F), |F \cap \mathcal{T}^1|, |F \cap \mathcal{T}^2|) = ((\ell(\partial F), \vec{v}(\partial F)), |F \cap \mathcal{T}^1|, |F \cap \mathcal{T}^2|).$$

We compare complexities lexicographically. Thus, the complexity of F is measured first by the complexity of its boundary, then by the *weight* of F, $\mathrm{wt}(F) = |F \cap \mathcal{T}^1|$, and then by the number of components of the intersections with the 2-skeleton of $\mathcal{T}$.

A normal surface is *least complexity* if it minimizes complexity among normal surfaces to which it is isotopic (but not necessarily normally isotopic).

A surface is *tight* if it minimizes complexity, ranging over *all* those surfaces to which it is isotopic.

A tight normal surface is clearly least complexity, and as a consequence of Proposition 7.1, a normal essential surface of least complexity is tight. But, this does not hold in general for surfaces that are not essential: for example, a normal boundary parallel annulus may be least complexity but after tightening no longer normal.

We first recover normalization of an essential surface. We will apply this with surfaces whose boundaries are tight, hence least length.

Proposition 7.1. *Suppose that X is a triangulated, irreducible manifold with incompressible boundary. If $F \subset X$ is a tight, properly embedded, essential surface, then F is normal.*

Proof. To prove F is normal we must show that it meets each tetrahedron Δ in a collection of disks whose boundaries are normal curves of length 3 or 4. We adopt the view taken in [BDTS12], showing F meets each tetrahedron in pieces that are incompressible and edge incompressible.

If any component of $F \cap \Delta$ is compressible in Δ, then, by an innermost disk argument, we obtain a compressing disk avoiding all other components of $F \cap \Delta$, and hence $F \cap \Delta$ is compressible inside Δ.

Because F is essential, the boundary of any compressing disk D for $F \cap \Delta$ is trivial in F. Because X is irreducible, compressing along D yields a surface F' that is isotopic to F, but for which either $|F \cap \mathcal{T}^1|$ or $|F \cap \mathcal{T}^2|$ has been reduced, a contradiction. It follows that $F \cap \Delta$ is the union of disks.

An *edge compressing disk* for a surface in Δ is an embedded disk E whose boundary $\partial E = e \cup f$, consists of two arcs, $e \subset \mathcal{T}^1$ and $f = E \cap F = \partial E \cap F$; see [BDTS12].

If some component of $F \cap \Delta$ has an edge compressing disk then, by an innermost disk argument, there is an edge compressing disk E for $F \cap \Delta$. If $e \subset \partial X$ then, because F is not boundary compressible, f is trivial in F. But compressing along E yields an isotopic surface F_0 (X is irreducible and has incompressible boundary) whose boundary length is reduced by at least two, contradicting the fact that $\partial F = \partial F_0$ is least length. And if e lies in an interior edge, then E can be used to guide an isotopy reducing $|F \cap \mathcal{T}^1|$, also a contradiction.

Then F meets each face in normal arcs. For otherwise, there is an arc whose ends both lie in the same edge, and an outermost such arc bounds an edge compressing disk. Then F meets the boundary of each tetrahedron in normal curves. And it is well known, see Thompson [Tho94], that if any such curve has length greater than 4 we see an edge compressing disk for F in the boundary of the tetrahedron. □

0-efficient triangulations. First we recall the definition of 0-efficient triangulations from [JR03]. A triangulation of a manifold X with nonempty boundary is 0-*efficient* if every normal disk is vertex-linking. (A normal disk is vertex-linking at vertex v if it consists of precisely one normal triangle from each tetrahedral corner meeting v.)

Moreover, if no boundary component of X is an S^2, then X does not contain any normal 2-spheres [JR03, Prop. 5.15]. In our setting, we use 0-efficient triangulations only in the situations without S^2 boundary components (since in the algorithm, we fill each such component with a ball). Note also that in the proposition below we can assume that X does not contain S^2 boundary components even if do not explicitly claim that X is obtained in an intermediate stage of the algorithm. Indeed, we assume that X is irreducible. Then an S^2 boundary component implies that X is a ball; however, the proposition also assumes that X contains an essential annulus or Möbius band.

We now establish the second result, that some summand of a non-fundamental incompressible annulus is an essential annulus. This applies to boundary compressible as well as essential annuli.

Proposition 7.2. *Let X be a triangulated, orientable, irreducible manifold with incompressible boundary and a 0-efficient triangulation. Let A be an incompressible annulus or Möbius band that has tight boundary and is least complexity and normal. Suppose that A can be written as a non-trivial sum $A = B + C$ where B is connected and $\partial B \neq \emptyset$. Then B is an essential annulus or Möbius band with tight boundary.*

7.1 Proof of Proposition 7.2

Sketch of the proof. Our proof is loosely modeled on Jaco and Tollefson's proof of [JT95, Th. 6.5]. Apart from using slightly different notion of complexity, we also have to add additional ingredients when A is a boundary parallel annulus.

As we will see, the core of the proof is to show that B is essential. For contradiction we assume that B is not essential. The first important step is to find out what are the possible

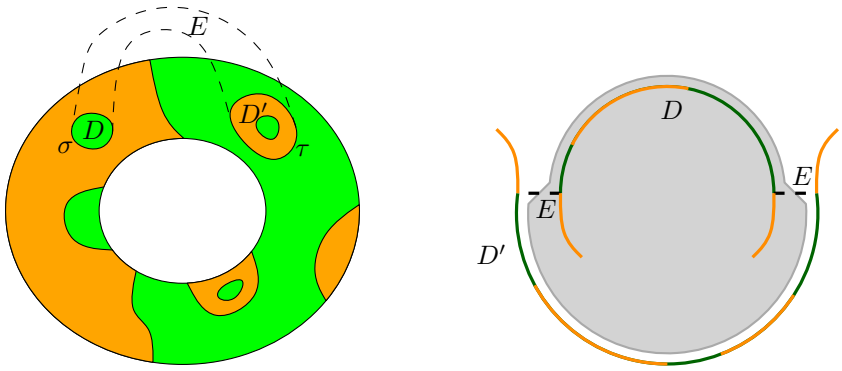

Figure 8: The annulus A separated by trace curves into patches (left); E is an exchange annulus. The right part shows a cross-section of $D \cup E \cup D'$ assuming that D and D' are not adjacent across E. Then $D \cup E \cup D'$ separates, after a slight isotopy, two components of A.

patches when A is decomposed by trace curves from the normal sum $A = B + C$; see Figure 8 left. If A is essential (annulus or Möbius band), then disk patches as well as half-disk patches can be ruled out following [JT95] (disk patches avoid ∂X whereas half-disk patches contain a single arc on ∂X); see Lemmas 7.8 and 7.9. After ruling out such patches we can deduce that every intersection curve is essential in B, that is a spanning arc or a core curve. This already mean that C intersects B in a very specific way and both cases can be ruled out along [JT95]; see Lemma 7.11.

If A is not essential, then A is a boundary parallel annulus by Proposition 4.2. In this case we do not know how to rule out disk patches but we still can rule out half-disk patches (Lemma 7.9); here we use that simplification of the boundary has higher priority than simplification of the interior in our notion of complexity. Since A is boundary parallel, there is an annulus $A_{\partial X}$ to which A is parallel and together they bound a solid torus T in X. Because there are no half-disk patches, we can show that one of the exchange rectangles for the sum $A = B + C$ is inside this torus and it meets A and $A_{\partial X}$ only in essential arcs. However, with such a rectangle A cannot be boundary parallel; see Lemma 7.10 for details. This finishes the sketch of the proof and now we provide the details.

Because A is incompressible, ∂A is essential by Proposition 4.2. Without loss of generality, we will assume that the sum $A = B + C$ lexicographically minimizes $(|\partial B \cap \partial C|, |B \cap C|)$, the number of boundary intersections and the total number of intersection curves, over pairs (B', C') where B' and C' are locally snug surfaces isotopic to B and C, respectively. Since $\partial A (= \partial B + \partial C)$ is assumed tight, we have, by Lemma 6.8, that ∂B and ∂C are tight, and because $|\partial B \cap \partial C|$ is minimized, snug.

Lemma 7.3. *Either the conclusion of Proposition 7.2 holds, or B is a boundary parallel annulus and every component of C is an incompressible annulus, Möbius band, torus or Klein bottle.*

Proof. No component of C has Euler characteristic $\chi > 0$: Because X is 0-efficient, no normal surface is a sphere, nor a projective plane, for then its normal double would be a normal sphere. And, also by 0-efficiency, any disk has boundary a trivial vertex linking curve that survives normal addition, and is present in ∂A—a contradiction.

Then every component has $\chi = 0$ and it is an annulus, Möbius band, torus, or Klein bottle. No component is a compressible annulus since these have a trivial boundary component (Proposition 4.2) and this contradicts the fact that both summands have essential boundary.

Since B is connected and $\partial B \neq \emptyset$, B is either an annulus or Möbius band. Also by Proposition 4.2, a Möbius band is essential and satisfies the conclusion of Proposition 7.2. So does an annulus, unless it is boundary compressible, and hence boundary parallel, by Proposition 4.2. □

We proceed with the proof of Proposition 7.2 under the assumption that B is a boundary parallel annulus.

When A is formed as the normal sum $B+C$, it is partitioned into patches coming from B and C, as was discussed in Section 5.1, and we have exchange surfaces attached to the curves separating the patches; see Figure 8 left.

It follows that no exchange surface is a Möbius band. As noted in Section 5.1, this occurs only when an intersection loop is one-sided in both summands.

Define a *half disk* to be a disk that is halfway properly embedded in X, that is, an embedded disk whose boundary meets ∂X in a single arc. Note that a boundary compressing disk for a surface is a half disk whose boundary meets the surface in the complementary arc, but the reverse does not hold in general, for the arc may not be essential in the surface.

An exchange rectangle or annulus E meets four patches of A. A pair D, D' of these patches are said to be *adjacent across* E if they meet opposite boundary curves σ and τ of E, but from the same side (we again refer to Figure 8 left).

Lemma 7.4. *σ bounds (half) disk in A if and only if τ both bounds a (half) disk in A.*

Proof. We prove that if σ bounds a (half) disk in A, then τ bounds a (half) disk in A. The reverse implication is proved by interchanging B and C and remarking that in this proof we do not use the extra assumptions on B.

The surface A is either essential or a boundary parallel annulus, and it is incompressible by the assumptions.

Suppose that σ bounds a disk D in A. Then $E \cup D$ is a disk which, after a slight isotopy, meets A only in τ. Since A is incompressible, τ bounds a disk in A as claimed.

The same argument works when A is essential and, say, σ bounds a half disk H. Then $E \cup H$ is a half disk which, after a slight isotopy, meets A only in τ. Since A is not boundary compressible, τ bounds a half disk in A.

We conclude by showing that τ bounds a half disk when σ bounds a half disk and A is a boundary parallel annulus (assumed to have tight boundary). To obtain a contradiction, suppose that σ bounds a half disk H but τ is an essential arc in A. Then $E \cup H$ is, after a slight isotopy, a boundary compressing disk meeting A in the arc τ. Since A is parallel to an annulus $A_{\partial X} \subset \partial X$, their union bounds a solid torus in X. The rectangle E is a disk properly embedded in this solid torus.

Indeed, if E is outside the solid torus, consider a boundary compressing disk D_A for A meeting A in τ inside the solid torus. Then the disk $D_A \cup E \cup H$ meets the core curve of $A_{\partial X}$ exactly once, implying that it is a non-separating disk and therefore a compressing disk for ∂X—a contradiction.

As soon as we know that E is inside the solid torus, we have that the boundary of E meets $A_{\partial X}$ in a pair of exchange arcs that each span $A_{\partial X}$ by Lemma 6.5, and meets the annulus A in one curve σ that is trivial in A and the other τ that is a spanning arc for A. Therefore,

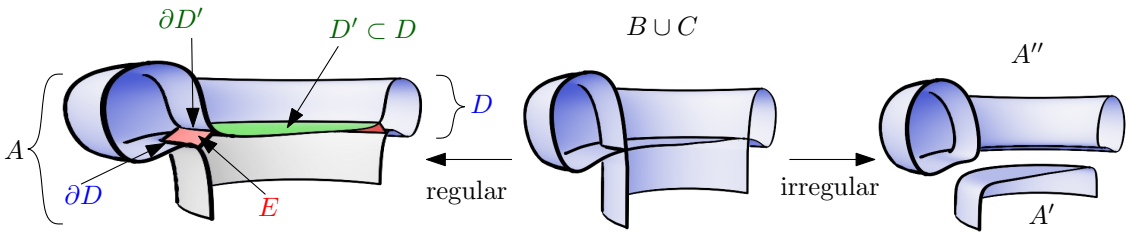

Figure 9: Non-adjacent disks, $D' \subset D$ (or half disks $H' \subset H$ by replacing D with H and D' with H'). The picture is actually drawn for half disks and the disk case is obtained by doubling.

when restricting to $A_{\partial X}$, we get that two corners of E are in one component of ∂A and the other two in the second one. When restricting to A, we get that three corners are in the same component and the other corner in the second—a contradiction. □

Lemma 7.5. *Suppose A is essential. If σ and τ bound disks in A, then they are adjacent across E.*

Proof. Suppose that σ and τ bound non-adjacent disks D and D'. They either are disjoint, or one is a sub-disk of the other, say $D' \subset D$.

When disjoint, the union $D \cup E \cup D'$ is a sphere that, after a slight isotopy, separates components of A (since spheres separate in irreducible manifolds)—a contradiction. See Figure 8 right.

Suppose then that $D' \subset D$; see Figure 9. Let A' be the surface obtained from A by removing D and replacing it with $E \cup D'$.

The union of the disk D and a slight offset of the disk $E \cup D'$ bounds a ball, across which the disks are isotopic (X is irreducible). So A', the result of this disk swap, is isotopic to A. Also note that performing an irregular exchange (fold) at this intersection loop produces a surface with two components: one is A', and the other, A'', is a torus obtained by identifying the ends of the annulus $D \setminus D'$.

Because E has zero weight, we have $\mathrm{wt}(A) = \mathrm{wt}(A') + \mathrm{wt}(A'')$. But, because this was not a regular exchange, $A' \cup A''$ is not normal, and there is an abnormal arc bounding a half disk in some face by Lemma 5.5. If this half disk meets ∂X, then $A' \cup A''$ either is boundary compressible or is not least length, both contradictions.

Thus, the half disk lies in the interior and can be used to guide an isotopy of $A' \cup A''$ that removes two intersections with the 1-skeleton. But this implies that $A' \cup A''$ can be isotoped to have strictly less weight than A. This is a contradiction since the component A' has lower complexity, but is isotopic to the tight surface A. □

Unfortunately, the above proof contradicts minimal interior weight and does not apply when A is a (non-essential) boundary parallel annulus, which may not be normal when tight. Fortunately, the half disk version contradicts tight boundary and can be applied when A is essential or a boundary parallel annulus.

Lemma 7.6. *If σ and τ both bound half disks in A, then they are adjacent.*

Proof. Suppose to the contrary that σ and τ bound half disks H and H' that are not adjacent across E. The half disks H and H' are either disjoint, or, say, $H' \subset H$.

If disjoint, then the union $H \cup E \cup H'$ is a properly embedded disk, that after a slight isotopy, separates components of A, which is a contradiction.

Now suppose that $H' \subset H$; see Figure 9. Replacing H with $E \cup H'$ is a disk swap across a ball that produces a surface A' isotopic to A. But notice that performing an irregular rather than regular switch at this intersection curve produces a surface with two components: one is A', and the other is an annulus A'' formed by identifying the ends of the rectangle $H \setminus H'$. The irregular switch on the intersection arc yields irregular switches at the endpoints which are intersections of the boundary curves.

So while $\ell(\partial A) = \ell(\partial A') + \ell(\partial A'')$, the curve $\partial A' \cup \partial A''$ is not normal, contains an abnormal arc by Lemma 5.5, and so there is an isotopy reducing its length. Since $\partial A' \cup \partial A''$ is isotopic to a curve of length strictly lower than ∂A, each of its components has length strictly lower than ∂A, contradicting the minimality of the length of ∂A. □

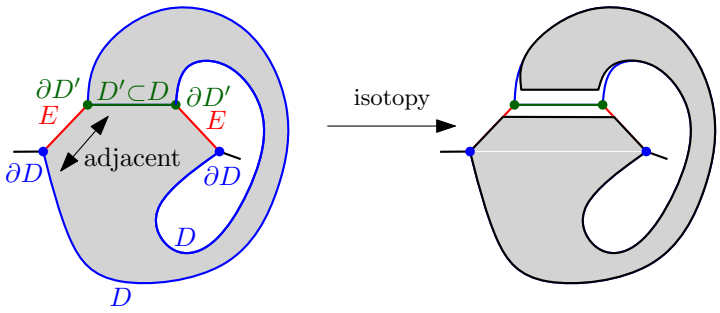

Figure 10: Adjacent disks; here $D' \subset D$. The union $D' \cup E \cup D$ bounds a ball after slight isotopy.

In the next two lemmas we will utilize the following observation:

Observation 7.7. *Suppose that D and D' are disks that are adjacent across the exchange annulus E. If the disks are disjoint, then $D \cup E \cup D'$ is, possibly after a slight isotopy, a sphere bounding a ball.*

Proof. If D and D' are disjoint, we get the observation immediately, since X is irreducible. If not, say if $D' \subset D$, then fix D and slightly isotope the interior of the disk $E \cup D'$ off D to the side of the exchange annulus; see Figure 10. After the isotopy $D \cup E \cup D'$ is a sphere bounding a ball. □

Similarly, if H and H' are half disks adjacent across the exchange rectangle E, then the union $H \cup E \cup H'$ together with ∂X bound a ball, possibly after a slight isotopy. Indeed, the union $H \cup E \cup H'$ is a properly embedded disk, after first perhaps slightly isotoping, say $E \cup H'$, when H and H' are not disjoint. Its boundary $\partial H \cap \partial E \cap \partial H'$, perhaps slightly isotoped, bounds a disk in ∂X and together these disks bound a ball in X.

We also describe how a surface A obtained as a normal sum $A = B + C$ can be obtained, under certain additional conditions, as a normal sum $A = B' + C'$ guided by some of the normal exchanges of $A = B + C$. Let $\mathcal{E}$ be the set of exchange bands for the normal sum

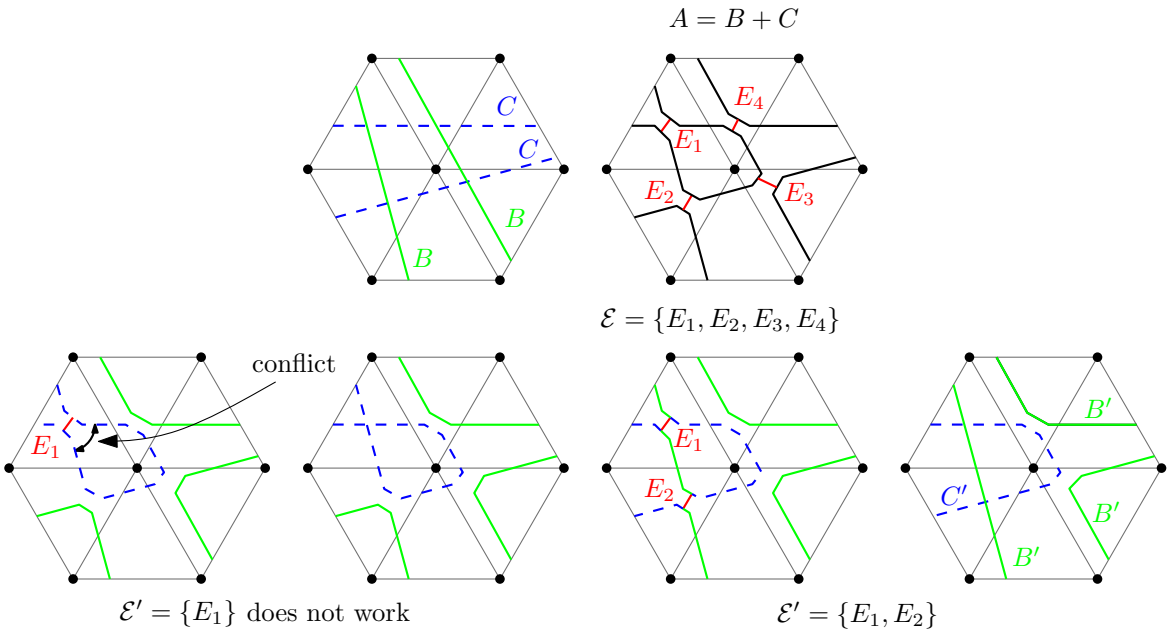

Figure 11: Summing B and C (upper left picture) yields A (upper right picture) with the corresponding set of exchange bands/arcs $\mathcal{E} = \{E_1, E_2, E_3, E_4\}$. Some subset $\mathcal{E}' \subset \mathcal{E}$ may not be consistent (two lower left pictures); however, if it is consistent (two lower right pictures), then summing B' and C' yields A again.

$A = B + C$. Note that $|\mathcal{E}| = |B \cap C|$. We say that a subset $\mathcal{E}' \subset \mathcal{E}$ is a *consistent subset* if the induced patches, components of $A \setminus \mathcal{E}'$, can be bicolored so that two patches have different colors if they either lie on opposite sides of the same trace curve, or, are adjacent across an exchange band. (If this happens for two sides of a same patch, then in particular it cannot be bicolored.) In this case, we can see that $\mathcal{E}'$ is the set of exchange bands for a normal sum $A = B' + C'$ where B' and C' are each the union of patches of a single color, connected across $\mathcal{E}'$. See Figure 11. The same analysis holds for subsets of exchange arcs for a curve sum $\alpha = \beta + \gamma$.

Lemma 7.8. *If A is essential, then there are no disk patches. Every intersection loop is essential in A, B, and C.*

This lemma does not apply to the boundary parallel case, because we cannot assume that disks are adjacent.

Proof. This argument appears in [JO84] and [JT95]. If there is a disk patch bounded by a trace curve, then by Lemmas 7.4 and 7.5, it is adjacent across an exchange annulus E to another disk. The adjacent disk is not a single patch,since if it were, $\mathcal{E} \setminus E$ would be a consistent subset, and we could expresses A as a sum $A = B' + C'$, where B' and C' are isotopic to B and C, respectively, but $|B' \cap C'| < |B \cap C|$, a contradiction. Thus, the adjacent disk contains trace loops. Pass to an innermost trace loop bounding a disk patch and repeat.

This process can be continued arbitrarily, and therefore must eventually repeat. Thus there is a shortest cycle of adjacent disks, each biting the tail patch of its predecessor across some exchange annulus; see Figure 12. Note that the cycle has length 1 when $D' \subset D$ for

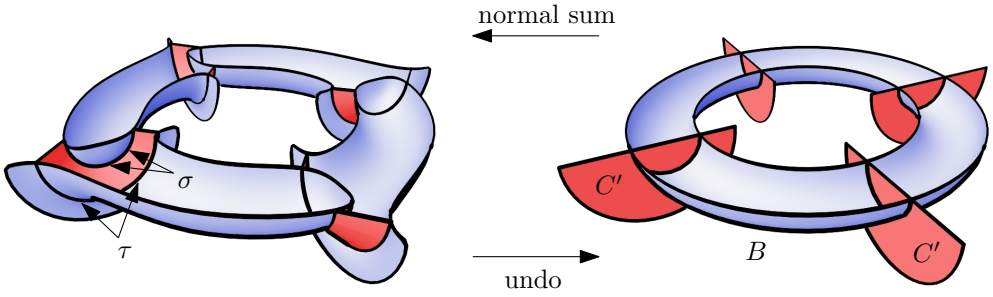

Figure 12: A cycle of adjacent half disks. Double the figure for a cycle of disks.

some pair of adjacent disks D and D'. By irreducibility, the union of a pair of adjacent disks and their exchange annulus is a sphere bounding a ball, and the union of these balls is a solid torus.

Let $\mathcal{E}'$ be the subset of $\mathcal{E}$ along which the cycle of disks is adjacent. Then $\mathcal{E}'$ is seen to be a consistent subset of $\mathcal{E}$ by coloring all annulus patches on the boundary of the solid torus with one color, and all other patches with the other. This expresses A as the sum of surfaces $A = B' + C'$ where B' is a normal torus bounding a solid torus.

Adding B' to C' corresponds to a fractional Dehn twist in B'. Since B' bounds a solid torus, there is an isotopy of the solid torus that undoes the twist and carries C' to A. So C' is isotopic to A but with strictly lower complexity $\mathrm{cpx}(C') = \mathrm{cpx}(A) - \mathrm{cpx}(B')$—a contradiction. □

We repeat the above argument to work in the context of trace arcs and half disks. This argument does apply to boundary parallel annuli, as their half disks are known to be adjacent. The proof will reach a contradiction to the tightness of ∂A.

Lemma 7.9. *No trace arc bounds a half disk in A.*

Proof. When A is a boundary parallel annulus, Lemma 7.8 does not apply, and trivial trace loops may be present. So we will refer to outermost (in A) half disks instead of half disk patches. An outermost half disk in A does not contain trace arcs, but, is not a patch if it contains trace loops. If there is an outermost half disk H bounded by a trace arc, then by Lemmas 7.4 and 7.6, it is adjacent across an exchange rectangle E to another half disk H'.

The adjacent half disk H' is not outermost: If H and H' are both outermost, then $H \cup E \cup H'$ is a properly embedded disk that meets ∂X in a disk bounding the curve $h \cup e \cup h' \cup e'$, where $h = H \cap \partial X$, $h' = H' \cap \partial X$ and e and e' are exchange arcs, the two components of $E \cap \partial X$. Since h and h' are sub-arcs of ∂B and ∂C, we can see, by undoing the exchanges e and e', that ∂B and ∂C bound a bigon and are thus not snug, a contradiction.

Thus, the adjacent half disk is not outermost and contains at least one trace arc. Pass to an innermost trace loop bounding a outermost half disk and repeat. This process can be continued arbitrarily, and therefore must eventually repeat. Thus, there is a shortest cycle of adjacent half disks, each biting the tail patch of its predecessor across some exchange

rectangle; see Figure 12. Note that the cycle has length 1 when $H' \subset H$ for some pair of adjacent half disks H and H'. By irreducibility and incompressibility of ∂X, the union of a pair of adjacent half disks and their exchange rectangle is a disk that co-bounds a ball with a disk in ∂X, and the union of all these balls is a solid torus meeting ∂X in an annulus.

Let $\mathcal{E}'$ be the subset of $\mathcal{E}$ along which the cycle of half disks is adjacent. Then $\mathcal{E}'$ is seen to be a consistent subset of $\mathcal{E}$ by coloring all rectangle patches on the boundary of the solid torus with one color, and all other patches with the other. This expresses A as the sum of surfaces $A = B' + C'$ where B' is a normal boundary parallel annulus.

Adding B' to C' corresponds to a fractional Dehn twist in B'. Since B' is boundary parallel, there is an isotopy of the solid torus that undoes the twist and carries C' to A. So C' is isotopic to A but with shorter length $\ell(C') = \ell(A) - \ell(B')$, a contradiction. □

We know that A is incompressible, and thus by Proposition 4.2 it is either an essential annulus or Möbius band, or a boundary parallel annulus. We deduce that in both cases it means that B is essential.

Lemma 7.10. *If A is a boundary parallel annulus, then B is essential (not a boundary parallel annulus).*

Proof. Suppose that B is a boundary parallel annulus. If $\partial C = \emptyset$ then $\partial B = \partial A$, and hence A and B are isotopic, contradicting the fact that A was chosen to have least complexity.

So we have $A = B + C$, where all three have non-empty tight boundary and A and B are both boundary parallel annuli. Then $\partial A = 2a$ and $\partial B = 2b$, where a and b are tight essential curves since A and B have tight boundaries. Since all normal coordinates of ∂A and ∂B are even, it follows that $\partial C = 2c$ for some tight essential curve c by Lemma 6.8.

Each pair, $2b$ and $2c$, of parallel curves bounds an annulus in ∂X. Normally isotope B and C so that these annuli are very thin and intersect in a collection of squares, each contained in a face of ∂X. Pick a particular square. Each of its corners is an endpoint of an intersection arc between B and C. All corners have the same normal sign. It follows that the exchange rectangles for corners on the same edge lie on opposite sides of A; see Figure 13.

Thus, there is an exchange rectangle R properly embedded inside the solid torus T that is bounded by A and $A_{\partial X} \subset \partial X$, the annulus into which A is isotopic. Then ∂R is a compressing disk for T since it meets A in a pair of arcs that are essential by Lemma 7.9 and meets $A_{\partial X}$ in a pair of arcs that are essential by Lemma 6.5. But this contradicts the fact that A is boundary parallel to $A_{\partial X}$. The unique, up to isotopy, compressing disk for T meets A and $A_{\partial X}$ each in a single essential arc. □

Lemma 7.11. *If A is essential then B is essential (not a boundary parallel annulus).*

Proof. We first note, as in [JT95], that each disk patch is incompressible and boundary incompressible. Any (boundary) compressing disk for a patch has its boundary in the essential surface A and therefore meets it in a trivial curve. This, in turn, implies the existence of a (half) disk patch.

In contradiction, suppose B is a boundary parallel annulus. By Lemma 7.8, B and C intersect in curves that are essential in both. These curves are either core loops or spanning arcs for B.

But they cannot be spanning arcs: If so, then there is a boundary compressing disk D for B so that $\partial D \cap B \cap \partial C = \emptyset$. Choose D to be such a boundary compressing disk that meets C

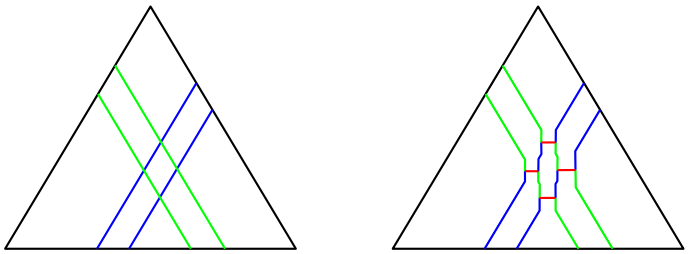

Figure 13: Annuli between $2b$ and $2c$ intersecting in a square. After regular exchanges there are four exchange arcs in which exchange rectangles start. At least one of these rectangles (actually two) have to be inside T.

in the minimal number of curves. Let D' be a disk bounded by an innermost loop, outermost arc, or when $C \cap D = \emptyset$, let $D = D'$. By minimality D' meets a patch in an essential curve and is thus a (boundary) compressing disk for a patch. This contradicts Observation 7.7.

We proceed assuming that $B \cap C$ is a collection of loops essential in both B and C. Then the patches of B and C are annuli because each component of C has zero Euler characteristic. As before, choose a boundary compressing disk D for B that meets C in the minimal number of curves.

By minimality of D and essentiality of patches, no curve of intersection is a closed loop or an arc with endpoints both in ∂X. It follows that γ, an outermost arc of intersection, co-bounds a bigon or half bigon D', with an arc $\beta \subset B \cap \partial D$. (Figure 3, with different labelling). By minimality of $|D \cap C|$, β is a spanning (essential) arc for an annulus patch $A_B \subset B$ and γ is a spanning arc for a patch $A_C \subset C$. The patches meet in either one (half bigon) or two (bigon) intersection loops. Moreover, because X is irreducible and has incompressible boundary, A_B and A_C are parallel across the solid torus that they bound. Thus, Figure 3 is a cross-section of the total intersection, *(half)bigon* $\times S^1$.

What does the normal addition $A = B + C$ do with the patches A_B and A_C? It cannot trade them as in the figure, because then, performing only the trade and no other exchanges produces normal surfaces B' and C' isotopic to B and C but with fewer intersections, contradicting our assumption that we had minimized $|B \cap C|$. Nor, in the half bigon case, can it attach A_B and A_C. For this means A has a boundary parallel annulus component. This rules out the half bigon case. And, in the bigon case, it cannot attach A_B to A_C along both curves, for if it did, a component of A would be a torus bounding a solid torus. So A_B and A_C are attached along one intersection loop, but not along the other. But then $A_B \cup A_C$ is parallel to the zero weight exchange annulus for the intersection loop where they were not attached.

Form a Haken sum A' by performing all regular exchanges, except perform an irregular exchange (fold) along the curve corresponding to the zero weight annulus. This is similar to the situation in Figure 9, although the context is a bit different. Then one component of A' is isotopic to A and the other, call it A'', is a torus bounding a solid torus. But $A' \cup A''$ is not normal because of the irregular switch. As in the end of the proof of Lemma 7.5, we conclude that $A' \cup A''$ is either boundary compressible, not least length, or, A is not tight. All are contradictions. □

This completes the proof of Proposition 7.2.

8 Constructing an annulus curve α

As usual, we assume that X is irreducible, orientable with incompressible boundary and presented via a 0-efficient triangulation.

Definition 8.1. *An* annulus curve α *is a properly embedded (multicomponent) normal curve in ∂X with the following property: There exists a collection $\mathcal{A}$ of pairwise disjoint properly embedded essential annuli in X such that $\alpha \subseteq \partial\mathcal{A}$ and α represents all normal isotopy classes of boundary components of $\mathcal{A}$ exactly once, i.e., for every annulus $A \in \mathcal{A}$ and every component γ of ∂A, there is exactly one component of α that is normally isotopic to γ.*

The following proposition provides an annulus curve that can be used to track essential annuli. Each boundary curve of a tight essential annulus either appears in the curve or meets the curve.

Proposition 8.2. *Let X be an irreducible, orientable manifold with incompressible boundary presented via a 0-efficient triangulation with t tetrahedra. Then there is a tight normal annulus curve α so that:*

(1) *α is maximal, by which we mean that if $A \subset X$ is an essential annulus or Möbius band whose boundary is tight and disjoint from α, then each boundary component of ∂A is normally isotopic to a component of α.*

(2) *$|\alpha|$, the number of components of α, is smaller than $4t$.*

(3) *$\ell(\alpha)$ is bounded by a computable function of t.*

The bound for $\ell(\alpha)$ we obtain from our proof is an $O(t)$-times iterated exponential, and this is currently a bottleneck of the whole algorithm.

The proposition follows in a simple way from the next two lemmas.

Lemma 8.3. *Suppose α_0 is a tight, normal annulus curve in the boundary of an irreducible, orientable manifold with incompressible boundary and a 0-efficient triangulation. Let A be an essential annulus with tight boundary that is disjoint from α_0, and such that a component of ∂A is not normally isotopic into α_0.*

Then there is a tight normal annulus curve α such that $\alpha_0 \subsetneq \alpha$ and $\ell(\alpha)$ is bounded by a computable function of t and $\ell(\alpha_0)$.

Proof. The curve α_0 can be regarded as a fence in the marked triangulation $(\mathcal{T}, \alpha_0 \cap \mathcal{T}^1)$. The annulus A has tight boundary disjoint from the fence α_0. By isotoping the interior of A (if necessary) while keeping its boundary ∂A fixed, we may assume that A is tight and hence, by Proposition 7.1, normal.

Write $A = F_1 + F_2 + \cdots + F_k$, a sum of connected fundamentals for the marked triangulation $(\mathcal{T}, \alpha_0 \cap \mathcal{T}^1)$. By Lemma 6.8 and by Proposition 6.10, respectively, each boundary ∂F_i is tight and disjoint from the fence α_0. Since $\ell(\partial F_i) \leq |F \cap \mathcal{T}^1|$, by Proposition 5.2, $\ell(\partial F_i)$ is bounded by a computable function of t and $\ell(\alpha_0)$, the number of marking points. By Proposition 7.2, each F_i with non-empty boundary is an essential annulus or Möbius band.

Moreover, there must be a summand with a boundary component that is not normally isotopic into α_0. Otherwise, each component of the boundary sum $\partial A = \partial F_1 + \partial F_2 + \cdots + \partial F_k$, would be normally isotopic to a component of the fence α_0. This would imply that the

summands are pairwise disjoint after a normal isotopy and, that each component of ∂A is itself normally parallel to a component of the fence α_0, contradicting our assumption.

Fix such a summand with a boundary component α' not normally isotopic into α_0. We will construct a new annulus curve α that contains both α_0 and α' (which ensures $\alpha_0 \subsetneq \alpha$).

Let A' be the summand if it is an essential annulus or twice the summand if it is a Möbius band.[13]

Let $\mathcal{A}_0$ be a collection of pairwise disjoint essential annuli witnessing that α_0 is an annulus curve, i.e., $\alpha_0 \subseteq \partial\mathcal{A}_0$ and α_0 represents all normal isotopy classes of boundary components of $\mathcal{A}_0$ exactly once. By construction of the annulus A', its boundary $\partial A'$ is disjoint from the fence α_0 and hence from $\partial\mathcal{A}_0$, and one boundary component of A' is the curve α'.

We isotope A', leaving the boundary fixed, to minimize components of intersection $|A' \cap \mathcal{A}_0|$, and we distinguish two cases.

The first case is that A' misses $\mathcal{A}_0$. Then $\mathcal{A} = \mathcal{A}_0 \cup A'$ is a collection of pairwise disjoint properly embedded essential annuli. We define α to be the annulus curve corresponding to $\mathcal{A}$. More precisely, α' is one of the boundary components of A'. If the other boundary component is normally isotopic to α' or to some component of α_0, we set $\alpha = \alpha_0 \cup \alpha'$, and otherwise, we set $\alpha = \alpha_0 \cup \partial A'$. Thus, α is an annulus curve, as witnessed by $\mathcal{A}$, $\alpha_0 \subsetneq \alpha$, and $\ell(\alpha)$ is bounded by a computable function of t and $\ell(\alpha_0)$ since $\ell(\partial A')$ is.

The second case is that A' intersects $\mathcal{A}_0$. In this case, since ∂A and $\partial\mathcal{A}_0$ are disjoint, a standard innermost loop argument shows that all curves of intersection are essential, i.e., core curves in the annuli; see the proof of Lemma 4.3.

If A' meets $\mathcal{A}_0$, let $\alpha'' \subset A'$ be the core intersection curve closest to α', and let $\tilde{A}$ be the corresponding annulus in $\mathcal{A}_0$. Then α' and α'' co-bound a sub-annulus $A'' \subset A'$ whose interior misses $\partial\mathcal{A}_0$.

Note that α'' splits the annulus component $\tilde{A}$ of $\mathcal{A}_0$ into annuli A_1 and A_2. Then α' is a boundary curve of both of the annuli $A'' \cup A_1$ and $A'' \cup A_2$, both of these are disjoint from $\mathcal{A}_0$, and at least one of them is essential.[14]

Suppose w.l.o.g. that $A''' = A'' \cup A_1$ is essential. Then $\mathcal{A} = \mathcal{A}_0 \cup A'''$ is a collection of pairwise disjoint essential annuli witnessing that $\alpha := \alpha_0 \cup \alpha'$ is an annulus curve (since one boundary component of A''' is α' and the other one is a boundary component of $\tilde{A} \in \mathcal{A}_0$ and hence already represented in α_0). Moreover, as in the first case, $\ell(\alpha)$ is bounded by a computable function of t and $\ell(\alpha_0)$ since $\ell(\partial A')$ and hence $\ell(\alpha')$ is. □

Lemma 8.4. *Let α be an essential curve embedded in a a connected closed orientable triangulated surface F with f faces, such that no pair of components of α are isotopic. Then $|\alpha| < f$.*

[13]The double of any normal surface F can be obtained by offsetting two copies, one to each side, of each normal disk of F. It follows that $2F$ is on the boundary of $F\widetilde{\times}[-1,1]$, an interval bundle over F. When the manifold is orientable, the bundle is a non-twisted product if and only if the surface is orientable. When the summand F is a Möbius band, $2F$ is the annulus $F\widetilde{\times}\{-1,1\}$. Since every curve in the boundary of an orientable manifold is two-sided, the boundary of the annulus is two copies of α'.

[14]First note that for each of these annuli, any core curve is parallel to α'', so a compression disk for either of these annuli would yield a compression disk for A' (and for $\tilde{A}$ as well), a contradiction. Thus, if the annuli are not essential, they both must be boundary compressible. Since we assume that ∂X is incompressible, it follows that each of the two annuli is boundary parallel, and thus each is of them co-bounds a solid-torus with an annulus in the boundary and is a longitudinal annulus in its respective solid torus (longitudinal means that it meets each meridional disk of the solid torus once). Then the union of the two solid tori is a solid torus in which $\tilde{A} = A_1 \cup A_2$ is longitudinal, demonstrating that $\tilde{A}$ is boundary parallel and contradicting the assumption that $\mathcal{A}_0$ consists of essential annuli.

Proof. The result holds trivially for a torus, which requires at least two faces to triangulate and allows α to have at most one component.

Now suppose that the genus of the surface is $g \geq 2$. We claim that then α has at most $3g-3$ components. We may assume that α is a maximal collection of non-parallel curves, and hence it decomposes F into p pairs of pants (spheres with 3 holes) [FM11, Sec. 8.3.1]. Each pair of pants has Euler characteristic -1, and because the Euler characteristic of the boundary of a pair of pants is zero, the Euler characteristic of the surface is additive over the pants, $\chi(F) = -p = 2-2g$. Because each curve is on the boundary of two pairs of pants, we have $|\alpha| = \frac{3}{2}p = 3g-3$.

We have $\chi(F) = 2-2g = f-e+v$, where $e = \frac{3}{2}f$ is the number of edges and v is the number of vertices. Then $g = \frac{f}{4}+1-v \leq \frac{f}{4}$ and $|\alpha| \leq 3g-3 \leq \frac{3}{4}f < f$, as desired. □

We now complete the proof of the main proposition.

Proof of Proposition 8.2. The annulus curve α can be constructed iteratively, starting with $\alpha = \emptyset$. If, at any stage, the maximality property (1) is not satisfied, we apply Lemma 8.3 to add a distinct component to α.

We claim that the process terminates after adding at most $4t$ components: By Lemma 8.4, each boundary component of X contains fewer components of α than it has faces (here, we are using that α is tight, so by Lemma 6.3, the fact that no two components of α are normally isotopic also implies that no two of them are isotopic). Thus, in total, α has fewer components than X has boundary faces and the number of boundary faces is bounded by $4t$. This is (2).

By Lemma 8.3, the first component added has length bounded by a computable function of t. Every subsequent component added has length bounded by a computable function of t and the total length of the preceding components. Since the number of components is less than $4t$, the total length of the curve is bounded by a computable function of t. □

9 Curves bounding boundary parallel annuli

In the previous section we constructed the annulus curve α, which will be used to bound the coefficients of essential annulus summands in the planar (almost) meridional surface P. In this section we construct Γ, a collection of curves bounding normal boundary parallel annuli. Later, the curves of Γ will act as fences, and will be used to rule out boundary parallel annulus summands altogether.

Proposition 9.1. *Suppose X is an irreducible, orientable manifold with incompressible boundary and presented via a 0-efficient triangulation with t tetrahedra. Let α be the tight normal annulus curve given by Proposition 8.2. Then there is a finite set Γ of tight essential curves, possibly mutually intersecting, such that:*

(1) *If B is a normal boundary parallel annulus with tight boundary that is disjoint from α, then each boundary component of B is normally isotopic either to a component of α or to a curve of Γ.*

(2) *$|\Gamma|$ and $\max_{\gamma\in\Gamma} \ell(\gamma)$ are bounded by a computable function of t.*

Proof. Let α be given by Proposition 8.2. Then α is a fence in the marked triangulation $(\mathcal{T}, \alpha \cap \mathcal{T}^1)$. Let Γ be the set of the boundaries of all boundary parallel annuli that are fundamental in the marked triangulation $(\mathcal{T}, \alpha \cap \mathcal{T}^1)$ and disjoint from the fence α. Then (2) follows from Proposition 5.2.

Now we want to verify (1). Let B be a normal boundary parallel annulus with tight boundary disjoint from the fence α. By isotoping B we can assume that B is least complexity.

If B is a fundamental, then its boundary has already been included in Γ and we are done. If not, then B can be written as a sum of fundamentals for $(\mathcal{T}, \alpha \cap \mathcal{T}^1)$, $B = F_1 + F_2 + \cdots + F_k$. By Proposition 7.2, each F_i with boundary is an essential annulus or Möbius band. Since B is disjoint from the fence α, each F_i has boundary disjoint from α. Hence, by Proposition 8.2, each F_i has boundary components normally isotopic to α. But as observed in the proof of Proposition 8.2, this implies that each boundary component of B is normally isotopic to α, as required. □

10 Planar meridional surfaces

In this section we consider a planar (almost) meridional surface P in X and a collection $\mathcal{A}$ of disjoint essential annuli. The collection $\partial\mathcal{A}$ of the boundaries of the annuli in $\mathcal{A}$ forms a collection of disjoint curves (loops) in ∂X, and ∂P is another collection of disjoint loops.

We want to move P by means of a self-homeomorphism $h\colon X \to X$ in such a way that the number of intersections of these two collections, $\partial\mathcal{A}$ and ∂P, becomes bounded; more precisely, we need a bound of the form $C(t)|\mathcal{A}| \cdot |\partial P|$. This is formulated in Proposition 10.3 below; the self-homeomorphism h is going to be one of two ways of changing the original embedding of X in S^3 in order to get a short meridian.

First we collect auxiliary results. We begin with a corollary of the main result of [MSTW13], which was developed for the purpose of proving a result in the spirit of Proposition 10.3.

Lemma 10.1 ([MSTW13, Cor. 1.6])**.** *Let S be connected surface, i.e., a connected compact 2-manifold with boundary, of genus g. Let $(\alpha_1, \dots, \alpha_m)$ be a system of disjoint curves (properly embedded arcs and loops) in S, and let $(\beta_1, \dots, \beta_n)$ be another such system. Then there is a homeomorphism $\varphi\colon S \to S$ fixing ∂S pointwise such that the total number of intersections of $\alpha_1, \dots, \alpha_m$ with $\varphi(\beta_1), \dots, \varphi(\beta_n)$ is at most $K(g)mn$, where $K(g)$ is a computable function depending only on g (in fact, $K(g) = O(g^4)$).*

We remark that for our further approach a bound of the form $K(g, m)n$ would also be sufficient. Such a bound, even independent of g, was obtained independently by Geelen, Huynh, and Richter [GHR13], but only under the additional assumption that the union of the β_i does not separate S. Thus, we cannot directly use their result here; the extra assumption could probably removed, but it is easier to use the bounds from [MSTW13].

We also need the following, probably standard, lemma.

Lemma 10.2. *Let $G = (V, E)$ be a graph with $n > 2$ vertices embedded in S^2, possibly with loops and multiple edges. Let us assume that no two parallel edges (connecting the same two vertices) and no two parallel loops (attached to the same vertex) are isotopic by an isotopy fixing the end-vertices and avoiding the other vertices. We also assume that there is no contractible loop ℓ; that is, both the interior and exterior of each loop contain a vertex. Then $|E| \le 3n - 6$.*

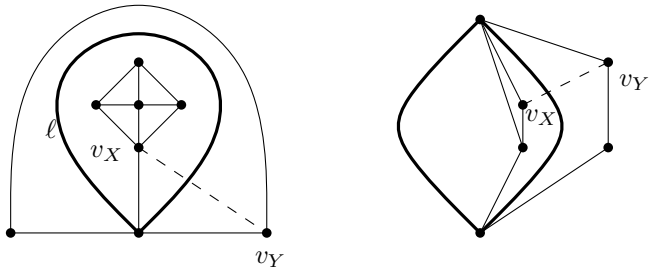

Figure 14: Removing loops and parallel edges.

Proof. If G contains neither loops nor multiple edges, then this is just the usual bound for the number of edges of a simple planar graph. It remains to resolve loops and multiple edges.

First let ℓ be a loop with an endpoint v. It splits S^2 into two regions X and Y. Let F_X resp. F_Y be the face of G inside X resp. Y bounded by ℓ. Then F_X has to contain a vertex $v_X \neq v$, for otherwise, ℓ can be contracted to v or it is isotopic to another loop with endpoint v.

Similarly, we have a vertex v_Y in F_Y. Note that v_X and v_Y are not connected with an edge. We can remove ℓ from the graph and connect v_X and v_Y with an edge, keeping the graph embedded in S^2 and satisfying the isotopy assumptions. This way we can remove all loops without increasing the number of edges; see Figure 14.

Similarly, if we have two parallel edges, we can remove one of them and add a new edge as compensation, reducing the number of pairs of parallel edges. In this way, we get a simple planar graph and the desired bound. □

Proposition 10.3. *Let X be an orientable, irreducible manifold with incompressible boundary. Let $P \subset X$ be a properly embedded planar surface that is either essential, or strongly irreducible and boundary strongly irreducible. Let $\mathcal{A} \subset X$ be a collection of pairwise disjoint essential annuli. Then there is a homeomorphism $h\colon X \to X$ so that $|\partial h(P) \cap \partial\mathcal{A}| < C|\mathcal{A}|\cdot|\partial P|$, where $C = C(t)$ is a computable function of the number t of tetrahedra in the triangulation of X.*

Proof. Using either Lemma 4.3 or Lemma 4.7, we may isotope P so that its intersection with $\mathcal{A}$ is essential, that is every component of $P \cap \mathcal{A}$ is a curve that is essential in both P and A. This implies that the result holds when P is a disk, for then P contains no essential curves, and thus $P \cap \mathcal{A}$ is empty. We proceed assuming $|\partial P| > 1$.

In $\mathcal{A}$ every intersection arc is a spanning arc and every intersection loop is the core curve of an annulus. This is illustrated below in the left picture, while the right picture shows the intersection curves in P:

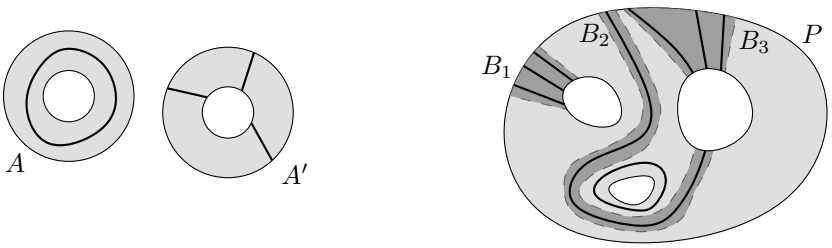

Say that two arcs belong to the same *parallel class* if they are isotopic in P. If $|\partial P| = 2$, then P is an annulus and there is at most one parallel class of intersection arcs. When $|\partial P| > 2$, form a planar graph by treating each boundary component as a vertex and each parallel class of arcs as an edge. The number of edges, hence parallel classes, is bounded by $3(|\partial P| - 2)$ by Lemma 10.2. We can cover the cases when P is an annulus or disk by reducing this last bound slightly. In all cases the number of parallel classes of arcs is bounded by $3(|\partial P| - 1) < 3|\partial P|$.

A *band* in X is an embedded, but not properly embedded, rectangle meeting ∂X in precisely its top and bottom sides. For each parallel class of intersection arcs in P, we may choose a band B_j that is a sub-surface of P, contains all intersection arcs in the class, and meets no other curves of intersection. Then $\mathcal{B}$, the union of all such bands, has at most $3|\partial P|$ components and contains all arcs, but no loops, of the intersection $P \cap \mathcal{A}$.

Next, let us draw the core curve α_i for every annulus in $A_i \in \mathcal{A}$, and a curve β_i parallel to the top and bottom sides (those in ∂X) in the middle of each band B_j. Let us think of these α_i and β_j as being (locally) horizontal and lying in the same level; then, again locally, A_i is a vertical "wall" through α_i and B_j is a vertical "wall" through β_j. We have the A_i and B_j fibered with segments, as in the left picture, and so the union $\mathcal{A} \cup \mathcal{B}$ has the structure of an I-bundle M_0 over $(\bigcup_i \alpha_i) \cup (\bigcup \beta_j)$, where I is the interval $[-1, 1]$; see the left picture below:

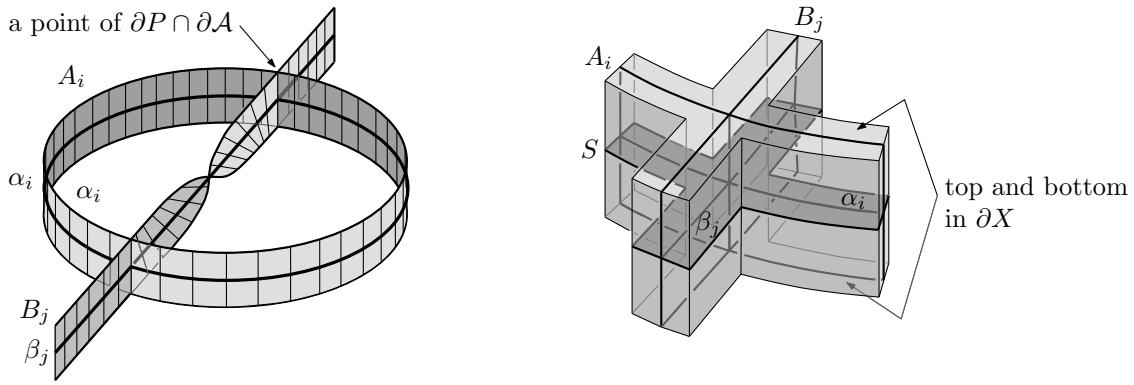

As the picture illustrates, some of the A_i or B_j may be twisted between the intersections with the others.

Next the I-bundle structure on $\mathcal{A} \cup \mathcal{B}$ can be extended to a sufficiently small regular neighborhood $N(\mathcal{A} \cup \mathcal{B})$. Indeed, we can consider the regular neighborhood as the star of $\mathcal{A} \cup \mathcal{B}$ (say in the second barycentric subdivision of some triangulation); see [RS72, Chapter 3]. Therefore, $N(\mathcal{A} \cup \mathcal{B})$ has locally structure as product of $\mathcal{A} \cup \mathcal{B}$ with I.

We obtain an I-bundle M over a base surface S forming a narrow ribbon along the α_i and the β_j. This is illustrated locally in the right picture above.

The plan is now to use Lemma 10.1 (untangling curves in a surface) for the systems of curves α_i and β_j within S, which yields a self-homeomorphism $\varphi\colon S \to S$ fixing ∂S pointwise, such that the number of intersections of the α_i with the $\varphi(\beta_j)$ is suitably bounded. Then we want to extend φ to a bundle self-homeomorphism $h\colon M \to M$ that is the identity over ∂S (i.e., on the vertical walls bounding M in the picture). After that, h can be extended identically to $X \setminus M$ and we will be done.

There are two issues to be handled. First, in order to use Lemma 10.1, bound the genus of each component of S by a computable function of t; we will actually obtain an $O(t)$ bound.

To this end, we observe that S is double-covered by a surface $\tilde{S} := N(\mathcal{A} \cup \mathcal{B}) \cap \partial X$. Let K be a component of S and $\tilde{K}$ be the corresponding double cover of K in $\tilde{S}$. In particular, $\chi(\tilde{K}) = 2\chi(K)$. For a surface F, we let $b(F)$ denote the number of boundary components, and define $g_e(F) := 2 - \chi(F) - b(F)$. If F is connected, this value is known as the Euler genus of F. Then we get

$$g_e(K) = 1 + g_e(\tilde{K})/2 - b(K) + b(\tilde{K})/2 \leq 1 + g_e(\tilde{K})/2.$$

Let $\tilde{Q}$ be a component of $\tilde{K}$ ($\tilde{K}$ has two components if K is orientable). Since $\tilde{Q} \subseteq \partial X$, we have $g_e(\tilde{Q}) \leq g_e(X^\partial_Q)$ where X^∂_Q is the component of ∂X containing $\tilde{Q}$. The Euler genus of X^∂_Q is bounded by $O(t)$, since $\partial X^\partial_Q = \emptyset$ and $|\chi(X^\partial_Q)| = O(t)$ (note that the number of triangles, edges and vertices in triangulation of X are all bounded by $O(t)$).

Altogether $g_e(\tilde{Q}) = O(t)$, and since $\tilde{K}$ has at most two components, $g_e(K) = O(t)$. Since the genus of a surface is at most twice the Euler genus, we also obtain $g(K) = O(t)$.

By applying Lemma 10.1 as announced above, working in each component K of S separately and then summing up, we obtain a self-homeomorphism φ of S, fixed pointwise on ∂S, such that the total number of intersections of the α_i with the $\varphi(\beta_j)$ is at most $C_0(t)|\mathcal{A}| \cdot |\mathcal{B}| \leq 3C_0(t)|\mathcal{A}| \cdot |\partial P|$, where $C_0(t)$ is a computable function of t.

It remains to deal with the second and last issue, namely, showing that φ extends to a bundle self-homeomorphism $h\colon M \to M$ that is identical over ∂S. Here we may assume w.l.o.g. that S, and hence M, are connected.

By the assumption, M is embedded in X, and so it is orientable. It follows that if the surface S is orientable, then M is actually the product $S \times [-1, 1]$, and the extension of φ to h is obvious.

So let S be non-orientable; then M is non-trivially twisted and there are no global coordinates. For a subset S' of S we will use a notation $S' \widetilde{\times} I$ for the subbundle of M consisting of points of M that project to S'. (In particular, we also regard M as $S \widetilde{\times} I$.)

For any connected non-orientable surface S there is a non-separating arc $\sigma \subset S$ with both endpoints on the same boundary component for which S_σ, which is S cut along σ, is an orientable surface. We also let $f_\sigma\colon S_\sigma \to S$ be the map gluing S_σ back to S.[15]

Thus, after cutting M along $\sigma \times I$, we obtain a product bundle M_σ, homeomorphic to $S_\sigma \times I$. In the boundary of M_σ we see two rectangle scars from cutting along $\sigma \times I$. We get the twisted bundle M back by gluing M_σ to itself along the rectangles so that the top of one is glued to the bottom of the other.

Now the given homeomorphism $\varphi\colon S \to S$ also takes σ to a curve $\varphi(\sigma)$ that has the same separation properties. We define $S_{\varphi(\sigma)}$ and $f_{\varphi(\sigma)}$ in the same way as S_σ and f_σ above. The homeomorphism φ also induces a homeomorphism $\varphi'\colon S_\sigma \to S_{\varphi(\sigma)}$ satisfying $f_{\varphi(\sigma)} \circ \varphi' = \varphi \circ f_\sigma$.

The homeomorphism φ' can be extended to a homeomorphism of the product bundles $h'\colon S_\sigma \times I \to S_{\varphi(\sigma)} \times I$ in two ways (by either keeping I or swapping it). By gluing back along the rectangular scars, h' induces a homeomorphism $h\colon M \to M$.

Recall that M is orientable since it embeds into X. Since we had two choices for h' we select one for which h is an orientation preserving automorphism. It follows that whenever K is a boundary component of S, then h is the identity on $K \widetilde{\times} I$ (it cannot flip I here since such a flip would reverse the orientation on ∂M, hence on M). □

[15] It is not difficult to find such an arc σ in the projective plane or the Klein bottle with a single hole. Any other nonorientable surface with nontrivial boundary can be obtained by adding handles (not across the desired arc σ) and holes to one of the two surfaces above (adding a handle increases the non-orientable genus by 2).

11 Proof of the short-meridian theorem

We already have almost all of the ingredients ready to prove Theorem 1.4, following the outline from Section 2.

We assume that X is irreducible, has incompressible boundary (which we may assume to be nonempty), embeds in S^3, and has a 0-efficient triangulation with t tetrahedra. Note that the second conclusion of the following lemma may require a re-embedding of X into S^3.

Lemma 11.1. *X contains a planar surface P so that:*

(1) *P is essential, or, strongly irreducible and boundary strongly irreducible, and*

(2) *∂P is meridional or almost meridional in some embedding of X in S^3.*

Proof. Since X embeds in S^3, we can apply the result of Fox [Fox48] that shows X may be embedded so that $S^3 \setminus interior(X)$ is a collection of handlebodies.

Then we may view X as the exterior, $X = S^3 \setminus N(\Gamma)$, where Γ is a graph consisting of a spine of each handlebody. In this context, we may apply Theorem 3 of Li [Li10] that states that X contains a planar surface that is either: (1) meridional, strongly irreducible, and boundary strongly irreducible, or (2) almost meridional and essential , or (3) non-separating, almost meridional, and incompressible. By Lemma 4.5, case (3) reduces to case (2), and the lemma follows. □

Lemma 11.2. *Let P be a surface satisfying the conditions of Lemma 11.1, and let $h\colon X \to X$ be a homeomorphism. Then $h(P)$ satisfies the conditions of Lemma 11.1 for some re-embedding of X.*

Proof. Because the homeomorphism h maps any disk in X to a disk in X, P is essential if and only if $h(P)$ is essential; and, P is strongly irreducible and boundary strongly irreducible if and only if $h(P)$ is strongly irreducible and boundary strongly irreducible.

Let $e\colon X \to S^3$ be the embedding for which P is (almost) meridional. Then $r := e \circ h^{-1}\colon X \to S^3$ is a re-embedding of X, and any component $e(\mu) \subset e(\partial P)$ bounds a disk in $S^3 \setminus X$ if and only if $r(h(\mu))$ bounds a disk in $S^3 \setminus (r \circ h(X))$. Then P is (almost) meridional in the original embedding if and only if $h(P)$ is meridional in the re-embedding. □

We can thus place additional constraints on P. Let α be the tight essential annulus curve given by Proposition 8.2 and let Γ be the set of tight essential curves bounding boundary parallel annuli given by Proposition 9.1.

Assumption 11.3. *Among planar surfaces P satisfying the conclusion of Lemma 11.1 choose P to minimize, in this order:*

(1) *$|\partial P \cap \alpha|$;*

(2) *$\mathrm{cpx}(\partial P)$, and hence $\ell(\partial P)$; and*

(3) *$|\partial P \cap (\alpha \cup \Gamma)|$.*

The next lemma shows that P's intersections with α are bounded by a linear function of $\chi(P)$ for t fixed.

Lemma 11.4. *Under Assumption 11.3, we have*

(1) $|\partial P \cap \alpha| \leq C_0(t)|\partial P|$, *where* $C_0(t)$ *depends on* t, *the number of tetrahedra;*

(2) ∂P *is tight and essential; and*

(3) $\{\partial P\} \cup \{\alpha\} \cup \Gamma$ *is pairwise snug.*

Proof. We have $\alpha \subset \partial\mathcal{A}$, where $\mathcal{A}$ is a collection of pairwise disjoint essential annuli. We choose $\mathcal{A}$ to minimize $|\mathcal{A}|$ subject to $\alpha \subset \partial\mathcal{A}$. Then $|\mathcal{A}| \leq |\alpha|$, because each $A \in \mathcal{A}$ must contribute at least one unique component to α.

By Proposition 10.3, there is a homeomorphism of X so that the image of the planar surface, call it P, satisfies $|\partial P \cap \partial\mathcal{A}| \leq C(t)|\mathcal{A}| \cdot |\partial P|$ for a suitable $C(t)$. Now Proposition 8.2 guarantees that $|\alpha| \leq 4t$, and hence $|\partial P \cap \partial\mathcal{A}| \leq 4C(t)t|\partial P|$ by Assumption 11.3.

Because P is either essential or strongly irreducible, ∂P consists of essential curves (Lemma 4.6). Thus, ∂P can be tightened; this may possibly increase $|P \cap \partial\mathcal{A}|$.

However, since α and Γ are tight, using Lemma 6.7 repeatedly, we can make $\{\partial P\} \cup \{\alpha\} \cup \Gamma$ pairwise snug within their normal isotopy classes. In particular, after this step $|P \cap \partial\mathcal{A}| = i(\partial P, \partial\mathcal{A})$ where $i(...)$ is the geometric intersection number. This again guarantees that $|P \cap \partial\mathcal{A}|$ is minimized.

Therefore, we can simultaneously achieve ∂P tight and $\{\partial P\} \cup \{\alpha\} \cup \Gamma$ pairwise snug. Hence both of these properties hold under Assumption 11.3. □

Lemma 11.5. *P can be isotoped, without changing ∂P, to be normal or almost normal.*

Proof. If P is essential, then, since ∂P is tight, P itself can be tightened without changing ∂P. Then P is normal by Proposition 7.1.

If P is strongly irreducible and boundary strongly irreducible, then the main result of [BDTS12] states that P is isotopic to an almost normal surface. Moreover, in the proof of Proposition 3.1 of [BDTS12] it is assumed that ∂P is least length (see Lemma 3.9), which is satisfied when ∂P is tight. The additional normalization steps taken there isotope the interior of P without changing its boundary, so ∂P is also fixed in the almost normal case. □

The average length argument. We mark the triangulation $\mathcal{T}$ of X with marking $M = (\alpha \cup \Gamma) \cap \mathcal{T}^1$. Thus, the (almost) meridional, (almost) normal planar surface P can be written as a sum of fundamental (almost) M-normal surfaces, $P = \sum k_i F_i$, and its boundary is the sum of the boundary curves of the fundamentals:

$$\partial P = \sum k_i \partial F_i.$$

Since ∂P is essential and tight, the boundary of each summand is essential and tight by Lemma 6.8. Each F_i falls into at least one of the following categories:

1. $\partial F_i = \emptyset$;
2. F_i is almost normal;
3. $\chi(F_i) > 0$ and F_i is normal;
4. $\chi(F_i) = 0$, F_i is normal, and one of the following hold:
   (a) F_i is a compressible annulus;

(b) F_i meets α;

(c) F_i is a boundary parallel annulus disjoint from α;

(d) F_i is an essential annulus or Möbius band disjoint from α.

5. $\chi(F_i) < 0$ and F_i is normal.

We will bound the total length $\ell(\partial P)$ by bounding the coefficients of each of these types in the boundary sum $\partial P = \sum k_i \partial F_i$. Obviously, we can ignore summands with empty boundary (Case 1). Since there can be at most one exceptional piece in the almost normal case, we have $k_i \leq 1$ in Case 2.

Lemma 11.6. *There are no normal summands with $\chi(F_i) > 0$.*

Proof. Such summands are either spheres, or projective planes, or disks. As we mentioned in Section 7, normal spheres contradict 0-efficiency of the triangulation of X.

Projective planes are excluded because X is irreducible. As for disks, since X has incompressible boundary, they also have trivial boundary. But this contradicts the fact that each summand has essential boundary. □

This excludes Case 3, and we proceed with Case 4a.

Lemma 11.7. *No summand is a compressible annulus.*

Proof. Because X has incompressible boundary, a compressible annulus has trivial boundary by Proposition 4.2. This would contradict the fact that every summand of ∂P is essential and tight by Lemma 6.8. □

The next lemma supplies a bound for Case 4b, although it does not need the assumptions of Case 4b in full strength.

Lemma 11.8. *$\sum k_i < C_0(t)|\partial P|$, where the sum is restricted to those F_i for which $\partial F_i \cap \alpha \neq \emptyset$.*

Proof. Addition of (almost) M-normal surfaces implies M-normal addition of their boundary curves which is additive with respect to intersections with the tight essential fence α by Proposition 6.10. Thus the sum of the coefficients k_i of fundamentals F_i meeting α is bounded by the total number of intersection with α, which in turn is bounded by Lemma 11.4. □

Case 4c can be excluded:

Lemma 11.9. *No F_i is a normal boundary parallel annulus disjoint from α.*

Proof. If not, then we can write $P = B + P'$, an (almost) M-normal sum, where $B = F_i$ is a boundary parallel annulus. All summands of P have tight boundary by Lemma 6.8, so $\partial B = 2b$ a pair of normally parallel tight essential curves. By the construction of Γ (Proposition 9.1), b is normally isotopic to either a component of α or an element of Γ, and therefore to a fence. By Proposition 6.10, each point of $\partial P' \cap b$ has the same normal sign. But since $\partial B = 2b$, all intersections in $\partial B \cap \partial P'$ have the same normal sign. This contradicts Lemma 5.6. □

Next, we want a bound for Case 4d.

Lemma 11.10. *$k_i < C_1(t)|\partial P|$ for each F_i that is a normal essential annulus or Möbius band disjoint from α, with a suitable $C_1(t)$.*

Proof. If F_i is a Möbius band, let $F_i' := 2F_i$ and $k_i' := \lfloor k_i/2 \rfloor$, and otherwise, let $F_i' := F_i, k_i' := k_i$. Then we write $P = P' + k_i'F_i'$, where F_i' is an essential annulus. We wish to show that $k_i' \leq |\partial P \cap \alpha| - 1$; then the result follows from Lemma 11.4.

So we proceed by contradiction, assuming $k_i' \geq |\partial P \cap \alpha|$. Let f_1 and f_2 be the components of $\partial F_i'$.

Proposition 8.2 guarantees that f_1 and f_2 are each normally parallel to a component of the fence α, and thus, by Proposition 6.10, $\partial P'$ meets each component f_j in points with the same normal sign.

Since intersection arcs join intersection points of opposite sign (Lemma 5.6), each arc component of $P' \cap F_i'$ meets both boundary components of F_i' and is thus a spanning arc of F_i'. There are $n = \frac{1}{2}|P' \cap \alpha|$ such spanning arcs of intersection, and $\partial P'$ meets, say, f_1 in n positive intersections and f_2 in n negative intersections. From the view of boundary curves $\partial P = \partial P' + k_i'f_1 + k_i'f_2$, where adding copies of f_1 and f_2 is a fractional Dehn twist (with fraction $\frac{k_i'}{n}$) in each of those curves.

We have assumed that $k_i' \geq |\partial P \cap \alpha| = |\partial P' \cap \alpha| = 2n$, so the fraction is greater than 1. Then $\partial P' + (k_i' - n)f_1 + (k_i' - n)f_2$ is homeomorphic to ∂P. Moreover, the homeomorphism can be extended over the annulus F_i' to a homeomorphism of X that is a Dehn twist in F_i'. But this homeomorphism takes P to a surface with shorter length and the same number of intersection with α, contradicting our choice in Assumption 11.3. (This is another place where we may change the embedding of X.) □

Finally, it is straightforward to bound those summands in Case 5.

Lemma 11.11. *We have $\sum k_i \leq -\chi(P) < |\partial P|$, where the sum is restricted to those F_i with $\chi(F_i) < 0$.*

Proof. We have observed that all summands of P have $\chi \leq 0$. Those with $\chi = 0$ do not contribute to $\chi(P)$, and so $\chi(P) = \sum k_i\chi(F_i)$ for the summands with $\chi \leq -1$. It follows that $\sum k_i \leq -\chi(P)$ for these summands. □

We are ready to bound the average length of a component of ∂P.

Lemma 11.12. *We have*

$$\ell(\partial P) \leq L(t)|\partial P|,$$

where L is a computable function of the number of tetrahedra t.

Proof. Recall that we wrote ∂P as a sum of boundaries of (almost) normal M-fundamentals for the marked triangulation $(\mathcal{T}, M)$, where $M = (\alpha \cup \Gamma) \cap \mathcal{T}^1$. Proposition 5.2 bounds both the weight of any fundamental solution and the total number of fundamental solutions by computable functions of t and the number of marked points $m = |M| = \ell(\alpha) + \ell(\Gamma)$. These lengths are bounded by computable functions of t by Propositions 8.2 and 9.1, respectively. Thus the weight of any M-fundamental solution and the total number of M-fundamental solutions are bounded by computable functions of t only.

As in the proof outline in Section 2, let $\ell_{\max} := \max\{\ell(\partial F_i)\}$, the maximum taken over all normal or almost normal M-fundamental surfaces F_i (in the marked triangulation of X).

Because the length of a surface's boundary is at most its weight, $\ell_{\max}$ is bounded by a computable function of t.

Because length is additive, we have

$$\ell(\partial P) = \sum k_i \ell(\partial F_i),$$

where the sum is restricted to surfaces F_i with non-empty boundary.

If F_i is one of the four types of fundamentals that contribute to $\ell(\partial P)$, then $k_i \leq C_2(t)|\partial P|$, where $C_2(t) = \max(C_0(t), C_1(t))$, by Lemmas 11.8, 11.10, and 11.11 (sometimes the bound is much better).

Since the number of distinct fundamentals is bounded by a computable function of t, call it $C_3(t)$, we have $\sum k_i \leq C_3(t)C_2(t)|\partial P|$ over all summands that contribute to ∂P. The total length is then bounded by $\ell_{\max} \cdot C_3(t)C_2(t)|\partial P|$ as the lemma claims. □

Theorem 1.4 now follows. Because P is meridional or almost meridional, at least $|\partial P| - 1$ of its boundary components are meridians (note that $|\partial P| > 1$ by Lemma 11.6). Hence the average length of a meridian is at most

$$\frac{\ell(\partial P)}{|\partial P| - 1} \leq \frac{L(t)|\partial P|}{|\partial P| - 1} = \frac{|\partial P|}{|\partial P| - 1} L(t) \leq 2L(t).$$

Unlike the children of Lake Wobegon, some meridian must be at most average, and hence its length is bounded by $2L(t)$, a computable function of t. This completes the proof of Theorem 1.4.

12 Embedding 3-dimensional complexes

In this section we prove Corollary 1.2: we provide an algorithm for $\text{EMBED}_{3\to3}$. It uses the algorithm for $\text{EMBED}_{2\to3}$, as well as an S^3 recognition algorithm and an algorithm for $\text{EMBED}_{2\to2}$.

Let K be a 3-complex for which we want to test embeddability in $\mathbb{R}^3$. We assume, w.l.o.g., that K is connected. The idea is to replace every 3-simplex of K by a suitable 2-dimensional structure so that an embedding of this 2-structure ensures the embeddability of the 3-simplex.

We call a vertex v of K a *cut vertex* if removing v from K disconnects K. We let $K' := (\operatorname{sd} K)^{(2)}$ to be the 2-skeleton of the barycentric subdivision of K (see the paragraph below the description of the algorithm). We will show that if K is connected and without cut vertices, then K embeds in $\mathbb{R}^3$ if and only if K' does. And we will also show that the assumption that K does not contain cut vertices is achievable.

Description of the algorithm (assuming K connected).

1. If K is homeomorphic to S^3 (which can be tested, as in the algorithm for $\text{EMBED}_{2\to3}$), return FALSE.

2. If there is a vertex whose link[16] is not embeddable in S^2, return FALSE. (The embeddability in S^2 can be tested using [GR79] and S^2 recognition, for example.)

[16] We recall that the *link* of a vertex v in a simplicial complex K consists of all simplices σ of K that do not contain v and such that σ together with v forms a simplex of K.

3. If K contains a cut vertex v, consider two connected induced subcomplexes K_1 and K_2 of K such that $K_1 \cup K_2 = K$ and $K_1 \cap K_2 = \{v\}$, $K_1, K_2 \neq K$. (Note that such K_1, K_2 exist: after removing v from K we can possibly obtain more than two components, but we can merge them into two groups.) Run the algorithm for K_1 and K_2 separately and return TRUE if and only if both K_1 and K_2 embed in $\mathbb{R}^3$.

4. Run the algorithm for $\mathrm{EMBED}_{2\to 3}$ with $K' := (\operatorname{sd} K)^{(2)}$ and return its answer.

Geometric realizations and the barycentric subdivision. In this section we need to carefully distinguish a simplicial complex K and its *geometric realization* $|K|$. (In this section we use $|\cdot|$ solely for geometric realizations, although earlier it meant the number of connected components.) Given a complex K, we denote its barycentric subdivision by $\operatorname{sd} K$. See the next picture for an example of barycentric subdivision, and e.g. [Mat03] or almost any textbook on algebraic topology for a detailed treatment of this notion.

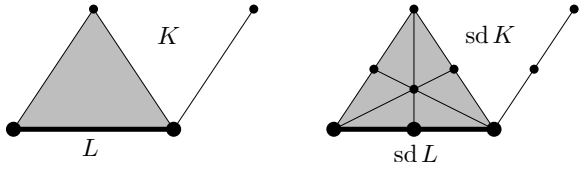

Given a subcomplex (or a face) L of K we also denote $\operatorname{sd} L$ the barycentric subdivision of L regarded as a subcomplex of $\operatorname{sd} K$. The geometric realizations of K and $\operatorname{sd} K$ can be canonically chosen so that $|K| = |\operatorname{sd} K|$ (and $|L| = |\operatorname{sd} L|$ for every subcomplex); we assume this canonical choice.

Correctness of the algorithm. Now we argue that the algorithm is correct modulo two lemmas proved below. In first step we exclude the case $K = S^3$ and thus, we can freely use that PL embeddability of K in $\mathbb{R}^3$ is equivalent to PL embeddability of K in S^3.

Further, if K PL embeds in $\mathbb{R}^3$ then the links of vertices PL embed in S^2,[17] so the answer in Step 2 is correct, and further we may assume that all the links embed in S^2.

The next lemma shows correctness of Step 3.

Lemma 12.1. *Let K be a connected simplicial complex such that link of each vertex embeds in S^2, and let K_1 and K_2 be two connected induced subcomplexes as in Step 3. Then K PL embeds in S^3 if and only if K_1 and K_2 PL embed in S^3.*

Finally, the correctness of Step 4 relies on the next lemma.

Lemma 12.2. *Let K be a connected simplicial complex without cut vertices, and let $K' = (\operatorname{sd} K)^{(2)}$. Then the following conditions are equivalent.*

(i) *K PL embeds in S^3;*

(ii) *K topologically embeds in S^3; and*

(iii) *K' topologically embeds in S^3.*

[17] Intersecting the PL embedding of K with a sufficiently small 2-sphere around the image of a vertex v, we get an embedding of the link of v in K.

Proofs of the lemmas. To finish the proof of correctness of the algorithm, it is sufficient to prove Lemmas 12.1 and 12.2. The proofs rely on the the 3-dimensional PL Schoenflies theorem (see, e.g., [Bin83, Theorem XIV.1]):

Theorem 12.3 (PL Schoenflies theorem for $\mathbb{R}^3$)**.** *Let $f\colon S^2 \to \mathbb{R}^3$ be a PL embedding.*[18] *Then there is a PL homeomorphism $h\colon \mathbb{R}^3 \to \mathbb{R}^3$ such that $h \circ i$ is a standard embedding of S^2 in $\mathbb{R}^3$ (as the boundary of a 3-simplex Δ^3). Moreover, h can be chosen to be the identity outside any given open set U that contains the bounded component of $\mathbb{R}^3 \setminus f(S^2)$. In particular, the bounded component of $\mathbb{R}^3 \setminus f(S^2)$ is a PL ball.*

This easily implies the following version for S^3 (which is also standard but we did not find a reference exactly in this setting):

Corollary 12.4 (PL Schoenflies for S^3)**.** *If $f\colon S^2 \to S^3$ is a PL embedding, then there is a PL homeomorphism $g\colon S^3 \to S^3$ such that $g \circ f$ is the standard inclusion of S^2 as the boundary of a hemisphere. In particular, the closures of both components of $S^3 \setminus f(S^2)$ are PL balls with boundary $f(S^2)$.*

Proof of Corollary 12.4. Choose a sufficiently fine PL triangulation of S^3 such that $f(S^2)$ avoids one of the closed d-simplices σ of S^3. By Newman's theorem [RS72, Cor. 3.13], the closure of the complement of a PL 3-ball in S^3 is a PL 3-ball, i.e., PL homeomorphic to a 3-simplex; in particular, $\overline{S^3 \setminus \sigma}$ is PL homeomorphic to a 3-simplex Δ^3_1.

Fix such such a PL homeomorphism $j\colon \overline{S^3 \setminus \sigma} \cong \Delta^3_1 \subseteq \mathbb{R}^3$. Then $j \circ f$ is a PL embedding of S^2 in $\mathbb{R}^3$. Thus, by Theorem 12.3, there is a PL homeomorphism h of $\mathbb{R}^3$ such that $h \circ j \circ f$ is a standard embedding of S^2 as the boundary of some 3-simplex Δ^3_2. Moreover, the PL homeomorphism $h \circ j$ witnesses that the closure of the component of $S^3 \setminus f(S^2)$ avoiding σ is a PL ball $B^3 \subset S^3$ with boundary $f(S^2)$. Furthermore, there is a PL homeomorphism k from Δ^3_2 to the closed lower hemisphere $H^3_- \subset S^3$ (e.g., with S^3 triangulated as the octahedral 3-sphere). By [RS72, Cor. 3.15], the PL homeomorphism $k \circ h \circ j\colon B^3 \cong H^3_-$ can be extended to a PL homeomorphism $g\colon S^3 \to S^3$, which has the desired property. □

Proof of Lemma 12.1. If K PL embeds in S^3, then both K_1 and K_2 PL embed in S^3 since they are subcomplexes of K. In sequel we assume that K_1 and K_2 PL embed in S^3 and we want to prove that K PL embeds in S^3.

The idea is very simple, we just want to transform an embedding of K_1 and K_2 so that the common vertex v protrude on the boundary of each and thus they can be joined together; see Figure 15 in one dimension less. It remains to show that such a transformation can be found.

From the assumptions on links of vertices and from $K_1 \neq K$ we deduce that the link of v in K_1 is planar (and thus different from S^2). Indeed, if it were homeomorphic to S^2, then the link of v in K would not embed in S^2 since the link of v in K_2 must be nonempty ($K_2 \neq \{v\}$ since $K_1 \neq K$). Similarly we can deduce that the link of v in K_2 is planar.

Let $f_1\colon |K_1| \to S^3$ be a PL embedding. By the previous observation we deduce that $f_1(v)$ is on the boundary of $f_1(|K_1|)$. Therefore, there is a geometric simplex σ in a small neighborhood of $f_1(v)$ such that $\sigma \cap f_1(|K_1|) = \{v\}$. Consequently, by the PL Schoenflies theorem (Corollary 12.4), there is a PL automorphism ψ of S^3 mapping $\partial\sigma$ to $S^2 \subset S^3$. In

[18] Formally, the standard PL model of S^d is the boundary $\partial\Delta^{d+1}$, and f is a PL map of the complex $\partial\Delta^3$ in $\mathbb{R}^3$.

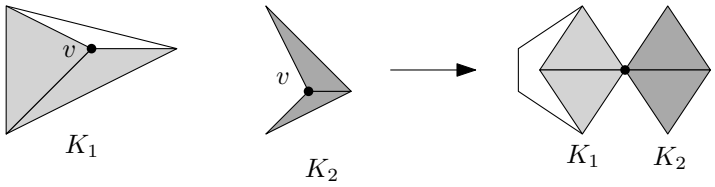

Figure 15: Transforming K_1 and K_2.

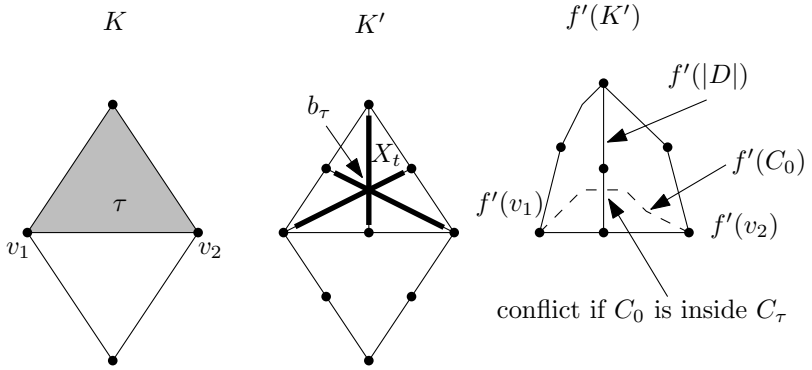

Figure 16: K, K' and $f(K')$ schematically drawn in the plane.

addition, we can assume that it maps the interior of σ to the upper hemisphere of S^3 and $f_1(v)$ to a pre-chosen point x on S^2. Altogether, $g_1 := \psi \circ f_1$ is a PL embedding mapping $|K_1|$ to the lower hemisphere of S^3 such that $g_1(|K_1|) \cap S^2 = \{x\}$.

Similarly, we can find a PL map $g_2 \colon |K_2| \to S^3$ such that $|K_2|$ is mapped to the upper hemisphere of S^3 and $g_2(|K_2|) \cap S^2 = \{x\}$. Finally, we can construct the desired PL embedding g of $|K|$ by setting $g(y) := g_1(y)$ if $y \in |K_1|$ and $g(y) := g_2(y)$ if $y \in |K_2|$. □

Proof of Lemma 12.2. Clearly (i)⇒(ii), and (ii)⇒(iii) since K' is a subcomplex of a subdivision of K. It remains to show (iii)⇒(i).

Since K' is 2-dimensional and since topological and PL embeddability coincide for embedding 2-complexes in S^3, there is an PL embedding $f' \colon |K'| \to S^3$. Let f_0 be the restriction of f' to $|K^{(2)}|$ (which is a subspace of $|K'|$). We want to extend f_0 to a PL embedding $f \colon |K| \to S^3$.

We will describe how to extend f_0 to each tetrahedron independently, and then we argue that these extensions can be done simultaneously, which yields the desired f. The argument is illustrated in Figure 16.

Let τ be a tetrahedron of K. By the PL Schoenflies theorem (Corollary 12.4), $f_0(\partial\tau)$ splits the sphere S^3 in two open components whose closures are PL homeomorphic to B^3. Let b_τ be the barycentre of τ and let C_τ be the component that contains $f'(b_\tau)$. We will argue that $f_0(|K^{(2)}|) \cap C_\tau = \emptyset$.

Recall that $f_0(\partial\tau)$ is disjoint from C_τ. For contradiction, let us assume that there is a component C_0 of $|K^{(2)}| \setminus \partial\tau$ such that $f_0(C_0) \cap C_\tau \neq \emptyset$, that is $f_0(C_0) \subseteq C_\tau$. Let $X_\tau := |(\operatorname{sd}\tau)^{(2)}| \setminus |(\operatorname{sd}\partial\tau)|$. We have $f'(X_\tau) \subseteq C_\tau$ since X_τ is connected. Let V_0 be the set of those vertices v of τ which belong to the closure of C_0. We have $|V_0| > 0$ since K is connected. Furthermore, $|V_0| > 1$ since K does not contain cut vertices. Therefore $|V_0| \geq 2$ and there

is a path P inside $|K^{(2)}|$ starting in a vertex v_1 of τ, ending in a vertex v_2 of τ, and with interior points in C_0.

Let D be the subcomplex of $(\operatorname{sd}\tau)^{(2)}$, homeomorphic to a disk, consisting of the simplices in the plane of symmetry of v_1 and v_2 (considering τ as a regular simplex); it passes through the other two vertices of τ and the midpoint of v_1v_2. By a double application of the PL Schoenflies theorem and using that the interior of $|D|$ belongs to X_τ (which maps to C_τ under f'), we have that $f'(|D|)$ splits the closure of C_τ into two parts, both PL homeomorphic to a ball. Since $f_0(v_1)$ and $f_0(v_2)$ are in different parts, $f'(|D|) \cap f_0(C_0) \neq \emptyset$, which is impossible since f' is an embedding.

We have deduced that $f_0(|K^{(2)}|) \cap C_\tau = \emptyset$ for each tetrahedron τ. By the PL Schoenflies theorem f_0 can be extended to a PL embedding of $K^{(2)} \cup \{\tau\}$ so that the interior of τ is mapped to C_τ. Moreover, if we consider two distinct tetrahedra $\tau_1, \tau_2 \in K$, then $C_{\tau_1} \cap C_{\tau_2} = \emptyset$. Indeed, if $C_{\tau_1} \cap C_{\tau_2} \neq \emptyset$, then $C_{\tau_1} \subseteq C_{\tau_2}$ because $\partial C_{\tau_2} = f_0(\partial\tau_2)$ is disjoint from C_{τ_1}. Similarly we deduce $C_{\tau_2} \subseteq C_{\tau_1}$ implying $\partial C_{\tau_1} = \partial C_{\tau_2}$ contradicting the assumption that τ_1 and τ_2 are distinct simplices.

Altogether, we can extend f_0 to every tetrahedron of K independently, obtaining the required PL embedding of K. □

Acknowledgment

We thank David Bachman, Ryan Derby-Talbot, and William Jaco for very helpful conversations. We further thank Tony Huynh for kind answers to our questions, Helena Nyklová for help with many pictures, and Isaac Mabillard for remarks on a preliminary write-up of the algorithm for $\text{EMBED}_{3\to3}$.

References

[AHT06] I. Agol, J. Hass, and W. Thurston. The computational complexity of knot genus and spanning area. *Trans. Am. Math. Soc.*, 358(9):3821–3850, 2006.

[Bac09] D. Bachman. Topological index theory for surfaces in 3-manifolds. Preprint arXiv:0901.0208, 2009.

[Bac13] D. Bachman. Stabilizing and destabilizing Heegaard splittings of sufficiently complicated 3-manifolds. *Mathematische Annalen*, 355(2):697–728, 2013.

[BDTS12] D. Bachman, R. Derby-Talbot, and E. Sedgwick. Almost normal surfaces with boundary. Preprint, arXiv:1203.4632, 2012.

[Bin59] R. H. Bing. An alternative proof that 3-manifolds can be triangulated. *Ann. Math. (2)*, 69:37–65, 1959.

[Bin83] R. H. Bing. *The geometric topology of* 3*-manifolds*, volume 40 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 1983.

[ČKM+12] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. Preprint, arXiv:1211.3093, 2012.

[ČKM+13] M. Čadek, M. Krčál, J. Matoušek, F. Sergeraert, L. Vokřínek, and U. Wagner. Computing all maps into a sphere. *J. ACM*, 2013. To appear. Preprint in arXiv:1105.6257. Extended abstract in *Proc. ACM–SIAM Symposium on Discrete Algorithms* (SODA 2012).

[ČKV13] M. Čadek, M. Krčál, and L. Vokřínek. Algorithmic solvability of the lifting-extension problem. Prepring, arXiv:1307.6444, 2013.

[FM11] B. Farb and D. Margalit. *A primer on mapping class groups*. Princeton University Press. Princeton, NJ, 2011.

[Fox48] R. H. Fox. On the imbedding of polyhedra in 3-space. *Ann. of Math. (2)*, 49:462–470, 1948.

[GHR13] J. Geelen, T. Huynh, and R. B. Richter. Explicit bounds for graph minors. Preprint, `http://arxiv.org/abs/1305.1451`, 05 2013.

[GL89] C. McA. Gordon and J. Luecke. Knots are determined by their complements. *J. Amer. Math. Soc.*, 2(2):371–415, 1989.

[GR79] J.L. Gross and R.H. Rosen. A linear-time planarity algorithm for 2-complexes. *J. Assoc. Comput. Mach.*, 20:611–617, 1979.

[Hak61] W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961.

[Hat82] A. E. Hatcher. On the boundary curves of incompressible surfaces. *Pacific J. Math.*, 99(2):373–377, 1982.

[Hem76] J. Hempel. 3-*Manifolds*. Princeton University Press, Princeton, N. J., 1976. Ann. of Math. Studies, No. 86.

[Hem79] G. Hemion. On the classification of homeomorphisms of 2-manifolds and the classification of 3-manifolds. *Acta Math.*, 142(1-2):123–155, 1979.

[HLP99] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999.

[Iva08] S. V. Ivanov. The computational complexity of basic decision problems in 3-dimensional topology. *Geom. Dedicata*, 131:1–26, 2008.

[Jac80] W. Jaco. *Lectures on three-manifold topology*, volume 43 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, R.I., 1980.

[JO84] W. Jaco and U. Oertel. An algorithm to decide if a 3-manifold is a Haken manifold. *Topology*, 23(2):195–209, 1984.

[JR03] W. Jaco and J. H. Rubinstein. 0-efficient triangulations of 3-manifolds. *J. Differential Geom.*, 65(1):61–168, 2003.

[JRS09] W. Jaco, J. H. Rubinstein, and E. Sedgwick. Finding planar surfaces in knot- and link-manifolds. *J. Knot Theory Ramifications*, 18(3):397–446, 2009.

[JS03] W. Jaco and E. Sedgwick. Decision problems in the space of Dehn fillings. *Topology*, 42(4):845–906, 2003.

[JT95] W. Jaco and J. L. Tollefson. Algorithms for the complete decomposition of a closed 3-manifold. *Illinois J. Math.*, 39(3):358–406, 1995.

[KMS13] M. Krčál, J. Matoušek, and F. Sergeraert. Polynomial-time homology for simplicial Eilenberg–MacLane spaces. *J. Foundat. of Comput. Mathematics*, 13:935–963, 2013. Preprint, arXiv:1201.6222.

[Kup11] G. Kuperberg. Knottedness is in NP, modulo GRH. Preprint, arXiv:1112.0845, 2011.

[Li10] Tao Li. Thin position and planar surfaces for graphs in the 3-sphere. *Proc. Amer. Math. Soc.*, 138(1):333–340, 2010.

[Mat97] S. V. Matveev. Classification of sufficiently large 3-manifolds. *Uspekhi Mat. Nauk*, 52(5(317)):147–174, 1997. Translation in Russian Math. Surveys 52 (1997), no. 5, 1029–1055.

[Mat03] J. Matoušek. *Using the Borsuk–Ulam theorem*. Springer, Berlin etc., 2003.

[MSTW13] J. Matoušek, E. Sedgwick, M. Tancer, and U. Wagner. Untangling two systems of noncrossing curves. Preprint, arXiv:1302.6475v3. Extended abstract in *Proc. Graph Drawing 2013*, 2013.

[MTW11] J. Matoušek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *J. Eur. Math. Soc.*, 13(2):259–295, 2011.

[Neu68] L. Neuwirth. An algorithm for the construction of 3-manifolds from 2-complexes. *Proc. Cambridge Philos. Soc.*, 64:603–613, 1968.

[Pap43] C. Papakyriakopoulos. A new proof for the invariance of the homology groups of a complex (in Greek). *Bull. Soc. Math. Grèce*, 22:1–154 (1946), 1943.

[RS72] C. P. Rourke and B. J. Sanderson. *Introduction to piecewise-linear topology*. Springer-Verlag, New York, 1972. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 69.

[Rub95] J. H. Rubinstein. An algorithm to recognize the 3-sphere. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, pages 601–611, Basel, 1995. Birkhäuser.

[Sch49] H. Schubert. Die eindeutige Zerlegbarkeit eines Knotens in Primknoten. *S.-B. Heidelberger Akad. Wiss. Math.-Nat. Kl.*, 3:57–104, 1949.

[Sch04] S. Schleimer. Sphere recognition lies in NP. Manuscript, available from `http://www.warwick.ac.uk/~masgar`, 2004.

[Sko95] A. B. Skopenkov. A generalization of Neuwirth's theorem on thickening 2-dimensional polyhedral. *Mathematical Notes*, 58(5):1244–1247, 1995.

[Sto00] M. Stocking. Almost normal surfaces in 3-manifolds. *Trans. Amer. Math. Soc.*, 352:171–207, 2000.

[Tho94] A. Thompson. Thin position and the recognition problem for S^3. *Math. Res. Lett.*, 1(5):613–630, 1994.

[Ton11] D. Tonkonog. Embedding 3-manifolds with boundary into closed 3-manifolds. *Topology and its Applications*, 158(9):1157–1162, 2011.